

# Learning Fourier-Constrained Diffusion Bridges for MRI Reconstruction

M. Usama Mirza, Onat Dalmaz, Hasan A. Bedel, Gokberk Elmas, Yilmaz Korkmaz, Alper Gungor, Salman UH Dar, Kader K. Oguz, and Tolga Çukur\* *Senior Member*

**Abstract**—Although MRI reconstruction requires a dealiasing transformation from undersampled to fully-sampled data, task-agnostic diffusion priors sample images via a denoising-based generative trajectory from an asymptotic start-point of Gaussian noise onto fully-sampled data. Since aliasing artifacts in MR images carry spatial structure deviating from Gaussian noise, this noise-governed trajectory can cause suboptimal artifact suppression. To address this limitation, we introduce the first Fourier-constrained diffusion bridge (FDB) for MRI reconstruction in the literature. Unlike task-agnostic diffusion priors, FDB does not rely on noise in its forward process and instead learns a dealiasing transformation between a start-point of undersampled data and the end-point of fully-sampled data. The start-point is derived via a stochastic Fourier-constrained degradation operator that removes a progressively growing set of spatial frequencies. Unlike cold/soft diffusion priors that use an asymptotic start-point of severely degraded measurements, FDB uses a realistically undersampled start-point to ensure closer alignment of model input between training and test distributions. Unlike existing diffusion bridges that use degradations based on weighted linear averages and noise addition, FDB implements degradations based on binary removal of compact k-space sets to conform to the physics of accelerated MRI. To further improve image quality, FDB leverages a novel sampling algorithm based on progressive dealiasing by continually correcting recovered k-space data across reverse diffusion steps. Demonstrations on brain MRI show that FDB outperforms competing methods by 4.5dB PSNR and 8.3% SSIM in within-domain and by 4.7dB PSNR and 16.4% SSIM in cross-domain reconstructions.

**Index Terms**—diffusion, bridge, generative, deep learning, MRI, reconstruction

## I. INTRODUCTION

MRI is a diagnostic powerhouse with exceptional soft-tissue contrast that suffers from long scan times. Acceleration via undersampled acquisitions helps lower operational costs and susceptibility to patient motion, albeit an ill-posed inverse problem must be solved to reconstruct images with the aid of image priors that improve problem conditioning [1], [2]. Given their high sensitivity, image priors based on deep learning have become pervasive in MRI reconstruction over the years [3]–[9]. A prevalent framework in this domain employs task-specific priors that are directly trained to perform

reconstruction, i.e., the priors capture a dealiasing transformation that maps undersampled to fully-sampled data [10]–[14]. The dealiasing transformation is typically mediated by a conditional model that receives a least-squares reconstruction of undersampled data as input, and that aims to output an image that would be consistent with fully-sampled data [15]–[27]. While these task-specific priors have been pervasive in the MRI reconstruction literature, they typically have limited representational capacity for detailed structural features [28]–[30], which in turn can compromise image fidelity [31]–[35].

*Task-agnostic diffusion priors:* An alternative framework that promises improved image fidelity trains task-agnostic priors to perform non-reconstruction tasks such as image generation [36], [37]. These priors use generative models to learn the distribution of fully-sampled data divorced from the MR imaging operator (i.e., decoupled from the influence of coil sensitivities and sampling patterns) [38]–[43]. Of particular promise, common diffusion priors employ a denoising-based generative trajectory to map pure Gaussian noise devoid of MRI signals as the asymptotic start-point onto fully-sampled data as the end-point (Fig. 1a). This transformation is mediated via a network-based recovery operator that receives a noisy image sample at a given time step and produces a less noisy sample as output [44]–[48]. To enable reconstruction of undersampled data, unconditional variants typically interleave denoising-based diffusion steps with data-consistency projections via the MR imaging operator during inference, which often provide only weak guidance toward the acquired data [49], [50]. In conditional variants, the recovery operator is instead trained to gradually recover less noisy samples starting from noise while also receiving undersampled measurements (e.g., zero-filled Fourier reconstruction of undersampled data) as a static input [51]–[55]. In both cases, the diffusion model acts as a generative prior within a posterior sampling framework, while the underlying generative trajectory remains governed by Gaussian-noise degradations rather than by the physics of k-space undersampling [49], [50].

Since the diffusion process in task-agnostic diffusion priors remains driven by Gaussian noise, MRI reconstruction involves inference dynamics that combine two distinct transformations [56], [57]: (i) the task-irrelevant denoising transformation induced by the reverse diffusion process, and (ii) the task-relevant dealiasing transformation induced by consistency with the acquired data. Because the structural characteristics of Gaussian noise and aliasing artifacts differ substantially for many undersampling patterns, these objectives often compete rather than complement one another [1]. In turn, inference

This study was supported in part by TUBA GEBIP 2015 and BAGEP 2017 fellowships, and by a TUBITAK 1001 Grant 121E488 awarded to T. Çukur (Corresponding author: Tolga Çukur, cukur@ee.bilkent.edu.tr).

Authors are with the Dept. of Electrical-Electronics Engineering and National Magnetic Resonance Research Center (UMRAM), Bilkent University, Ankara, Turkey, 06800.

can converge to solutions that only partially satisfy either data-consistency or fidelity to the clean MR image manifold, limiting reconstruction quality in task-agnostic diffusion priors [58]–[61].

*Cold/soft diffusion priors:* To improve performance in solution of inverse problems by increasing task relevance, several recent computer vision studies have considered cold or soft diffusion priors [62], [63]. Commonly, these methods train a prior to progressively map images degraded due to a known imaging operator (e.g., a large blur kernel) onto clean images reflecting ground truth. To derive intermediate samples during forward diffusion, cold diffusion priors granularize degradations from the imaging operator into multiple steps (e.g., a set of compact blur kernels), whereas soft diffusion priors use hybrid operators that combine such degradation with additive Gaussian noise [64]. However, similar to task-agnostic diffusion priors, cold/soft diffusion priors are formulated as general generative priors that recover images from heavily-degraded and hence weakly-informative start-points (Fig. 1b; 100-fold undersampled data for cold diffusion, pure Gaussian noise for soft diffusion). This results in a notable mismatch between the degradation level assumed by the prior (i.e., the distribution of images at the near-information-depleted start-point) versus the degradation level of actual MRI measurements (i.e., the distribution of measured images). Note that accelerated MRI exams yield diminishing returns beyond moderate undersampling rates, with sharply declining image quality and saturation in scan-time savings [65]. Consequently, many studies focus on moderate accelerations (e.g.,  $R=2-16$ ) to attain a practical balance between scan efficiency and image fidelity [66]. As cold diffusion priors are optimized for image recovery from heavily degraded inputs (e.g.,  $R>100$ ), their effectiveness is inherently constrained when reconstructing images from data acquired at these moderate ranges commonly used in MRI. As such, cold/soft diffusion priors can face elevated difficulty in reverse diffusion sampling [67], and suffer from suboptimal performance [62], [63]. Indeed, a recent study on MRI reconstruction suggests that cold diffusion priors may not offer benefits in image quality over task-agnostic diffusion priors [68].

*Diffusion bridges:* Diffusion bridges (DB) offer an alternative approach to improve task relevance by employing a finite start-point of realistically degraded images as their model input [69]–[72]. With degradation levels closely mimicking those prescribed for actual measurements, this approach delimits gross variability in the distribution of model inputs between training and test sets [69]. Existing diffusion bridges perform forward diffusion via linear-averaging degradation operators, which express intermediate image samples in the diffusion process as a noise-added linear average of degraded and clean images weighted according to time step [71], [72] (Fig. 1c). However, for MRI reconstruction tasks, additive noise in intermediate samples would force the recovery operator to partly learn a denoising transformation that restricts the task relevance of the diffusion prior. Moreover, linear averaging of undersampled and fully-sampled data would result in non-zero intensity across the entire k-space for all intermediate image samples, which is discrepant with the nature of accelerated

MRI acquisitions that exercise binary selection of removed k-space points [1]. The recovery operator must then learn to gradually amplify the entire set of unacquired k-space points at each reverse diffusion step, which might be prone to sampling inaccuracies. Thus, existing diffusion bridges might have limited utility for performant MRI reconstruction.

*Proposed method:* Here, we introduce a novel Fourier-constrained diffusion bridge (FDB) for improved performance in MRI reconstruction (Fig. 2). To enhance task relevance of the diffusion prior, FDB avoids noise corruption in its forward process to instead learn a dealiasing transformation that maps between undersampled and fully-sampled data. *Unlike task-agnostic or cold/soft diffusion priors based on an asymptotic start-point of severely degraded images*, FDB is formulated as a diffusion bridge that explicitly models conditional transport between moderately undersampled and fully-sampled MRI data, rather than synthesizing images from a near-information-depleted boundary condition. The realistic level of undersampling at the start-point ensures a rudimentary alignment between the dealiasing transformations performed by the recovery operator and required for MRI reconstruction. *While soft diffusion priors rely on noise addition*, FDB solely leverages undersampling degradations to enhance relevance for MRI reconstruction. To devise a diffusion process that is consistent with the physics of k-space undersampling in accelerated MRI, FDB leverages a novel Fourier-constrained degradation operator that removes a distinct set of spatial frequencies per forward diffusion step, thus gradually increasing the undersampling level across timesteps. In the absence of noise addition, the removed frequencies are selected randomly across timesteps and across the training set to attain a degree of stochasticity, which is postulated to be essential for diffusion modeling [73]. *Unlike conventional diffusion bridges based on linear-averaging degradation operators*, FDB's degradation operator attains improved task relevance by respecting the binary k-space point selection in accelerated MRI acquisitions and avoiding additive noise; and its recovery operator imputes a compact set of k-space points at each reverse diffusion step to facilitate reverse diffusion sampling. To further improve reconstruction performance during inference, *we introduce a novel sampling algorithm for FDB based on a learned correction term*. Unlike generic predictor-corrector schemes, the proposed algorithm enables progressive dealiasing across reverse diffusion steps to continually update estimates of recovered k-space points. Comprehensive demonstrations on brain MRI datasets indicate the superior performance of FDB against previous non-diffusion and diffusion priors. Code for FDB is available at <https://github.com/icon-lab/FDB>.

### Contributions:

- To our knowledge, FDB is the first diffusion bridge for accelerated MRI reconstruction in the literature.
- FDB leverages a novel stochastic, Fourier-constrained degradation operator based on random spatial-frequency removal to map the end-point of fully-sampled data onto the finite start-point of realistically undersampled data.
- For enhanced reconstruction, FDB leverages a novel sampling algorithm with a learned correction term to

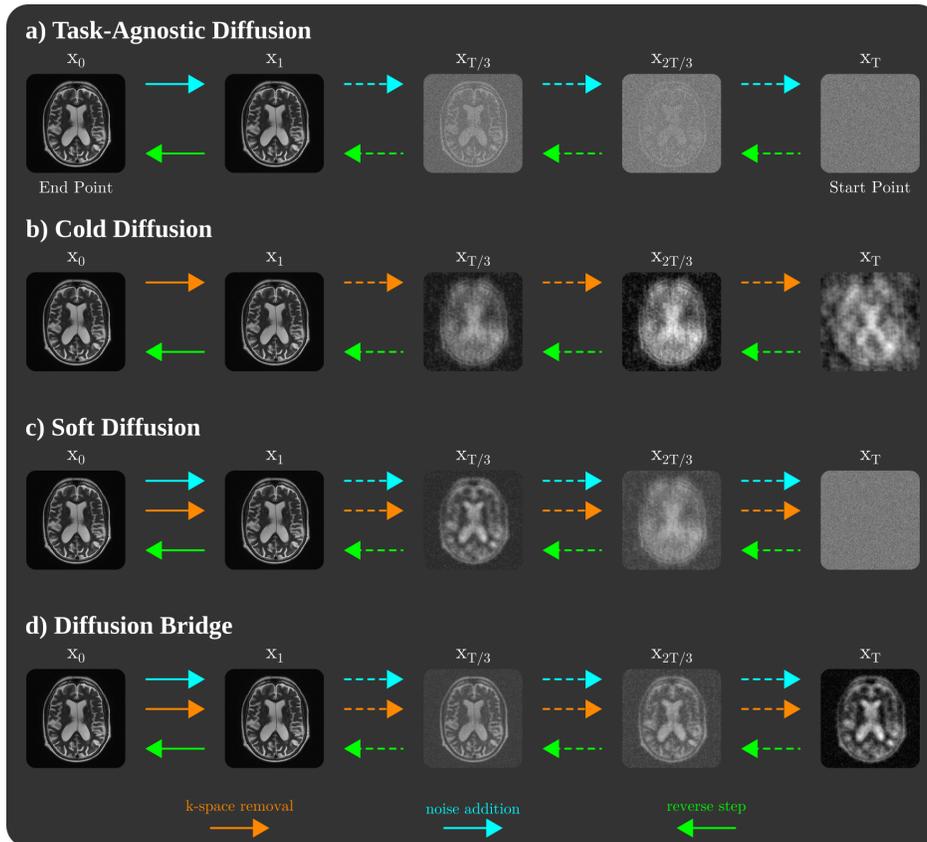


Fig. 1: Illustration of diffusion processes employed by **a)** task-agnostic diffusion priors, **b)** cold diffusion priors, **c)** soft diffusion priors, and **d)** conventional diffusion bridges.  $x_0$  is the clean image at the end-point, and  $x_T$  is the degraded image at the start-point.

achieve progressive dealiasing across reverse steps.

- Analytical derivations are provided for the optimal scheduling of the proposed correction term across reverse steps to maximize reconstruction performance.

## II. RELATED WORK

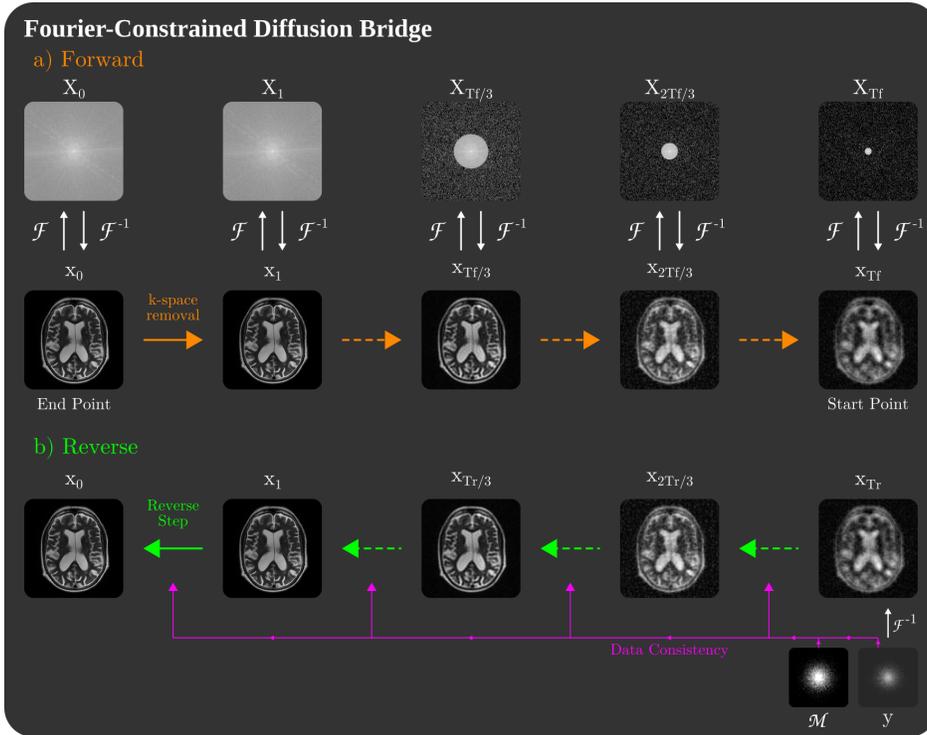
### A. Task-agnostic Diffusion Priors

Task-agnostic diffusion priors show high representational capacity for MR images, enabled via a stochastic process that gradually transforms between Gaussian noise and fully-sampled data [44], [49], [50], [73]. While promising results have been reported in MRI reconstruction [47], [51], [52], such priors sample the asymptotic start-point of the diffusion process from a white Gaussian noise distribution, and produce images via a denoising-based generative trajectory [73]. However, since this trajectory is governed by noise statistics rather than aliasing structure and can be weakly aligned with the dealiasing-driven trajectory required for MRI reconstruction, task-agnostic diffusion priors may suffer from limited task relevance [59], [60].

Several recent studies have focused on improving the effectiveness of task-agnostic diffusion priors in MRI reconstruction. For additional guidance, the recovery operator in conditional variants progressively denoises the pure Gaussian noise image at the start-point while receiving the undersampled measurements as a separate, static input [51], [52]. However, as the diffusion process remains governed by Gaussian-noise

degradations, inference must still reconcile denoising transformations induced by reverse diffusion steps with dealiasing transformations induced by measurement guidance. Since the characteristics of Gaussian noise and aliasing artifacts differ for many practical undersampling patterns, these two transformations can diverge across time steps in the diffusion process, resulting in under- or over-emphasis of conditional guidance and suboptimal recovery performance [69], [74]. In contrast, FDB employs a degradation operator that explicitly removes spatial frequency components without introducing noise corruption. This design ensures that the recovery operator performs targeted dealiasing rather than balancing competing transformations, enabling FDB to learn a prior directly optimized for MRI reconstruction.

Constraining the denoising diffusion process to high frequencies, HFS-SDE first high-pass filters white Gaussian noise and then adds this noise to clean MR images [46]. Importantly, this process does not selectively discard any frequency components of the MR images; rather, it superimposes high-frequency noise across the entire image while preserving the underlying image content. As a result, the forward degradations produce a noise-contaminated image, which differs from practical accelerated MRI acquisitions that generate aliasing-corrupted images. Consequently, HFS-SDE's starting point does not directly reflect the physical acquisition process of undersampled MR data. In contrast, FDB avoids additive noise altogether, and instead employs realistic undersampling degradations that induce variable-density sampling across k-space. Note that FDB's degradation strategy elicits a start-



**Fig. 2:** **a)** FDB’s degradation operator maps the clean image ( $x_0 = \mathcal{F}^{-1}(X_0)$  where  $X_0$  denotes the Fourier transform of  $x_0$ ) onto the finite start-point if the least-squares reconstruction of  $R'$ -fold undersampled data ( $x_{T_f} = \mathcal{F}^{-1}(X_{T_f})$  where  $X_{T_f}$  denotes the Fourier transform of  $x_{T_f}$ ). The recovery operator performs purely task-relevant dealiasing transformations. **b)** To reconstruct an  $R$ -fold undersampled acquisition  $y$ , sampling is initiated with  $x_{T_r} = A^H y$ , where  $A^H$  denotes the Hermitian adjoint of the imaging operator  $A$ , and  $T_r = \lfloor T_f \frac{(R-1)R'}{(R'-1)R} \rfloor$ .

point that is an aliasing-corrupted image with a k-space sampling density more closely aligned with accelerated MRI acquisitions, and this alignment facilitates image recovery during reverse diffusion.

### B. Cold Diffusion Priors

Another study has proposed a cold diffusion prior, CDiffMR, for single-coil MRI reconstruction [68]. While both FDB and CDiffMR employ undersampling-based degradation operators, key differences exist in the design of their diffusion processes. First, as a cold diffusion prior, CDiffMR is formulated as a generative prior that synthesizes images from a heavily-degraded start-point (i.e.,  $\approx 100$ -fold undersampled data), yielding a weakly-informative boundary condition for its diffusion process. At inference, the diffusion process is initialized at an intermediate timestep corresponding to the prescribed undersampling rate (e.g.,  $R=4$ ). Since CDiffMR’s prior is trained for recovery from substantially more degraded inputs, its effectiveness can be limited when recovery relies on a short temporal segment of the diffusion process [68]. In contrast, FDB is formulated as a diffusion bridge between two informative boundary conditions due to its finite start-point at a realistic degradation level, minimizing mismatch between training and inference, enabling more effective temporal use of the diffusion process, and thereby facilitating image recovery during reverse diffusion sampling.

Second, CDiffMR applies 1D undersampling with unordered removal of entire k-space lines, whereas FDB employs 2D undersampling with a gradually scheduled peripheral-to-central ordering for k-space point removal. This design choice

directly impacts the granularization of the degradation path. To achieve the same nominal acceleration, 1D masks have to eliminate a greater portion of low-frequency components, discarding k-space data that carry higher signal energy and are essential for stable recovery [1]. As a result, each diffusion step in CDiffMR removes a relatively large fraction of k-space support, yielding a coarse degradation trajectory that is challenging to invert. By contrast, FDB’s 2D peripheral-to-central masks progressively prune high-frequency components while retaining denser coverage of low frequencies throughout the diffusion process. This enables a finer-grained degradation path, in which each reverse diffusion step corresponds to a more tractable recovery problem. As a result, FDB’s start-point preserves a larger share of the underlying image energy available in fully-sampled data, providing more informative initial conditions and more stable recovery during reverse diffusion.

Finally, unlike generic correction loops that ignore mask evolution, the CCS algorithm is explicitly coupled to FDB’s non-revisiting frequency-domain degradation path, using the evolution of sampling masks and k-space energy across timesteps for selective refinement of imputed k-space points.

### C. Soft Diffusion Priors

An alternative for solving inverse problems in computer vision is soft diffusion priors, which generalize score-matching to accommodate hybrid operators that combine known deterministic degradations (e.g., blur, downsampling) with additive Gaussian noise [63]. While promising for natural image

restoration, existing soft diffusion formulations do not explicitly prescribe how to schedule k-space undersampling degradations across timesteps during forward diffusion, nor how to perform error correction on imputed frequency components during reverse diffusion in the context of accelerated MRI. Moreover, soft diffusion is formulated as a general generative prior that synthesizes images starting from an information-depleted boundary condition of pure Gaussian noise. When adopted for MRI reconstruction, a soft diffusion prior would simultaneously add noise and apply k-space undersampling in its forward process, causing the recovery operator to learn a joint denoising–dealiasing transformation with limited task specificity. Furthermore, soft diffusion employs aggressive noise degradation levels such that the start-point converges onto pure Gaussian noise devoid of MRI signals, as in task-agnostic diffusion priors. In turn, this mismatch causes the data distribution to diverge notably from that of undersampled MRI data, limiting the ability of soft diffusion to improve task relevance for MRI reconstruction.

To enhance reconstruction performance, the proposed FDB model is instead formulated as a diffusion bridge between moderately undersampled and fully-sampled data, employing a noise-free, frequency-domain degradation operator that explicitly represents MRI undersampling. This ensures that the recovery operator directly learns a dealiasing transformation and that the diffusion process remains aligned with the physics of MRI reconstruction. FDB also prescribes a peripheral-to-central ordering for k-space point removal and realistic degradation levels at the start-point of the diffusion process, together producing a physically meaningful progression of degradations that mimics the energy distribution of MRI spectra and ensures closer alignment between training and test distributions. Moreover, FDB introduces the CCS algorithm, explicitly coupled to its non-revisiting frequency-domain degradation path, enabling selective refinement of previously imputed k-space points guided by the evolution of sampling masks and k-space energy levels, beyond what generic predictor-corrector loops can achieve. This restores the iterative refinement behavior innate to diffusion models in the context of undersampling degradations. Collectively, these design elements render FDB better attuned to the physics and inference demands of MRI reconstruction.

#### D. Diffusion Bridges

Another emergent alternative in computer vision is diffusion bridges that learn transformations between two arbitrary distributions [69]. Conventional DBs derive intermediate samples via linear averaging of start- and end-points, followed by Gaussian noise addition [71], [72]. This design limits task relevance by forcing the bridge to learn an auxiliary denoising transformation. Moreover, when adopted for MRI reconstruction, linear averaging across the diffusion process would force the bridge to recover the entire set of k-space points at each step, elevating sampling inaccuracies. In contrast, FDB’s recovery operator learns to impute a compact set of k-space points at each reverse step, significantly facilitating diffusion sampling compared to conventional DBs.

### III. THEORY

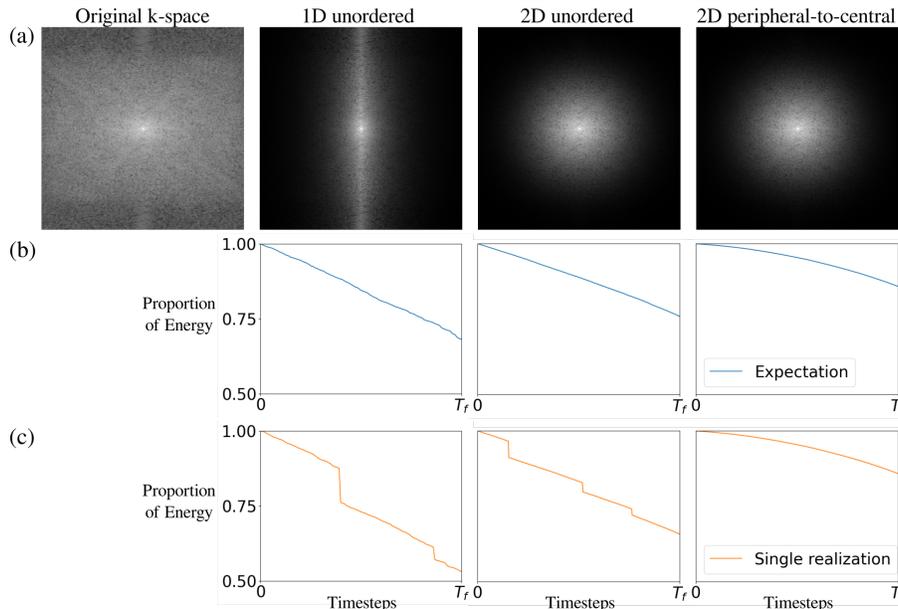
#### A. Fourier-Constrained Diffusion Bridges

MRI reconstruction aims to recover a high-quality MR image  $x$  given undersampled k-space acquisition  $y$ , linked via an imaging operator  $A=M\mathcal{F}(C)$  ( $M$ : sampling pattern,  $\mathcal{F}$ : Fourier encoding,  $C$ : coil sensitivities). To learn this task, FDB constructs a novel diffusion process based on binary removal-recovery of spatial-frequency components in MR images, mapping between a start-point  $x_{T_f}$  at timestep  $T_f$  taken as a coil-combined image undersampled at a rate  $R'$ , and an end-point  $x_0 = C^H \mathcal{F}^{-1}(y_{full})$  ( $H$ : Hermitian adjoint) taken as a coil-combined image derived from fully-sampled acquisitions (Fig. 2). FDB’s degradation operator exercises binary removal of a compact set of frequency components in each step, as this facilitates reverse diffusion sampling via gradual imputation of a small set of unacquired k-space points. FDB further utilizes a realistic degradation level  $R'$  to attain a closer alignment between the distribution of the model input in training versus test sets, and thereby improve learning of its recovery operator.

In this section, the formulation of FDB’s degradation and recovery operators are detailed. FDB is an image-domain diffusion model that expresses the intermediate samples across time steps in image domain. Yet, in the forward direction, the degradation operator is implemented by visiting Fourier domain for frequency removal via multiplication with binary matrices, and returning to image domain to obtain respective image samples. In the reverse direction, the recovery operator is implemented exclusively in image domain to avoid potential biases due to divergent energy levels at low versus high frequency components in k-space [1]. We derive the training objective of FDB’s recovery operator based on the interpretation that its forward steps are equivalent to a generalized diffusion process with linear degradations in image domain [63]. For the first time in literature, we introduce a continually-corrected sampling algorithm that incorporates a learned correction term to mediate progressive image dealiasing across reverse steps, and we present an analytical derivation of the optimal scheduling for its weight. Finally, we outline the inference procedures for MRI reconstruction on test data.

**A.1 Novel diffusion process:** The proposed FDB framework introduces a novel diffusion process that is uniquely different from existing methods: rather than relying on additive noise or asymptotic start-points, it directly models realistic k-space undersampling in MRI. A key design choice in FDB is the use of 2D undersampling with a peripheral-to-central schedule for k-space point removal during the forward diffusion process. This schedule simultaneously induces a variable-density sampling distribution that decreases toward the periphery of k-space and enforces a progressive ordering of point removal from higher to lower frequencies across diffusion steps.

MR images exhibit a Fourier-domain energy distribution that is heavily concentrated in low frequencies. Undersampling strategies that remove entire k-space lines, such as 1D masks, inevitably discard a relatively large number of these low-frequency coefficients to achieve a given acceleration factor. The resulting loss of high-energy components amplifies uncertainty in the recovery problem. FDB addresses this limitation



**Fig. 3:** Comparison of alternative k-space undersampling strategies for the forward diffusion process of FDB. **(a)** The log scale intensity distribution of an example MR image in k-space (Original k-space), and its masked versions with the expected sampling densities based on 1D unordered, 2D unordered, and the proposed 2D peripheral-to-central sampling strategies. **(b)** Expectation of the proportion of k-space energy remaining in intermediate image samples across the diffusion process for 1D unordered, 2D unordered and 2D peripheral-to-central sampling. **(c)** A single realization of the proportion of energy remaining in intermediate image samples across the diffusion process.

through 2D masks that can flexibly remove individual k-space points. By enforcing a peripheral-to-central removal order, higher-frequency coefficients are pruned first, while low-to-moderate frequencies—critical for accurate reconstruction—are preserved. Furthermore, since 2D masks can be configured to discard fewer k-space points per step (e.g., a minimum of  $n = 1$ ) than 1D masks (e.g., a minimum of  $n = 256$  for a readout with 256 k-space points), they induce a more gradual degradation over successive timesteps, allowing for finer discretization of the reverse process and improved stability during recovery. This combined design maintains denser coverage of informative frequency bands throughout the diffusion process, yielding a start-point that retains a substantially larger proportion of the energy from the fully-sampled acquisition. The result is a more informative and well-conditioned initialization for reverse diffusion. Note that this design can be viewed in analogy to noise variance scheduling in conventional diffusion priors, where noise degradations are scheduled to monotonically increase across timesteps, allowing the model to gradually learn to recover images from increasingly degraded inputs [44], [75].

To illustrate the motivations for our proposed design, Fig. 3a compares expected k-space coverage for a representative MR image under 1D unordered, 2D unordered, and 2D peripheral-to-central sampling at equal undersampling rates. These coverages were obtained by multiplying the original k-space data with the expected sampling densities under each strategy. In general, 2D strategies provide isotropically improved coverage of low-to-moderate frequencies compared with 1D sampling. Meanwhile, Fig. 3b shows the expected timecourse of k-space energy retained across forward diffusion steps, summarizing the average behavior of each sampling strategy. Although the expected sampling density is technically identical for 2D unordered and 2D peripheral-to-central masks,

2D unordered tends to discard a greater portion of lower-frequency components earlier in the diffusion process, which notably reduces the remaining energy at  $t = T_f$ . Consequently, 2D peripheral-to-central ordering achieves the highest energy retention consistently across  $0 < t \leq T_f$ , providing the most favorable starting condition for reverse diffusion in terms of the information available about the clean MR image. Fig. 3c depicts a single realization of the timecourse, illustrating the variability that arises from the random sampling masks that would be employed in practice. Note that 1D and 2D unordered masks both produce abrupt and irregular drops in retained energy levels, which can destabilize recovery. In contrast, peripheral-to-central ordering ensures smooth and gradual evolution of energy levels across timesteps, thereby stabilizing the reverse diffusion process. Together, these findings indicate that FDB’s 2D peripheral-to-central sampling provides a more informative initialization and a more stable trajectory for diffusion-based recovery.

**Degradation operator.** For a two-dimensional Cartesian k-space grid with  $N^2$  frequency components (where  $N$  is the number of grid points along each dimension), let  $k \in [1 N^2]$  denote linear component index,  $r(k) \in [0 r_{\max}]$ , and  $\phi(k) \in [0 2\pi)$  denote the k-space radius and angle of the  $k$ th component. In a forward step, the degradation operator performs frequency removal in Fourier domain via multiplication with a diagonal frequency-removal matrix  $\Lambda_t \in \mathbb{R}^{N^2 \times N^2}$ . At step  $t$ ,  $\Lambda_t(k, k)|_{k \in S_t} = 0$  for  $n$  components in the randomly selected set  $S_t$ , and  $\Lambda_t(k, k)|_{k \notin S_t} = 1$ . To improve the prior via monotonically increasing degradation across time [70], we observe that components selected at  $t$  should not overlap with previously selected components at  $\{t - 1, \dots, 1\}$ . To minimize the loss of energy as k-space points are removed, we further encourage a peripheral-to-central ordering of component selection, consistent with the

energy distribution of MR data in k-space [1]. We enforce these selection guidelines via a spatio-temporal point process over k-space and  $t$ :

$$S_t = \left\{ (k_1, \dots, k_n) : k_i \sim U[1, N^2] \text{ for } i = (1, \dots, n), \right. \\ \left. k_i \notin \bigcup_{\tau=1}^{t-1} S_\tau, r(k_i) > \bar{r}_t \right\}, \quad (1)$$

where  $U$  is uniform distribution,  $r_{max}$  is the maximum k-space radius, and  $\bar{r}_t$  is a radius threshold scheduled as:

$$\bar{r}_t = r_{max} - r_{max}(1 - \sqrt{R'}) (t/T_f). \quad (2)$$

Monotonically lowering  $\bar{r}_t$  to promote peripheral-to-central ordering towards the start-point, this scheduling ensures a degradation level of  $R'$  at  $t = T_f$  assuming that the number of components removed per step is  $n = \lfloor N^2(R' - 1)/(R'T_f) \rfloor$ .

Once random removal masks are determined according to the point process in Eq. 1, the relationship between  $X_t$  and  $X_0$  in Fourier domain is given as:

$$X_t = \left( \prod_{\tau=1}^t \Lambda_\tau \right) X_0 = \bar{\Lambda}_t X_0, \quad (3)$$

where  $\bar{\Lambda}_t \in \mathbb{R}^{N^2 \times N^2}$  is the cumulative frequency removal matrix at  $t$  (i.e., the product of removal matrices for  $\{t, \dots, 1\}$ ), and  $X_t, X_0 \in \mathbb{C}^{N^2 \times 1}$  are column-vectorized forms of respective complex Fourier-domain data. Note that the image-domain equivalent of the relationship in Eq. 3 can be expressed as:

$$x_t = \mathcal{F}^{-1}(\bar{\Lambda}_t X_0) = \mathcal{F}^{-1}(\text{diag}(\bar{\Lambda}_t) \otimes X_0), \quad (4)$$

$$= \mathcal{F}^{-1}(\text{diag}(\bar{\Lambda}_t) \circledast \mathcal{F}^{-1}(X_0)), \quad (5)$$

$$= \bar{\kappa}_t \circledast x_0. \quad (6)$$

In Eq. 4,  $x_t = \mathcal{F}^{-1}(X_t)$  denotes the image sample at  $t$ ,  $\text{diag}(\cdot)$  extracts the diagonal entries of  $\bar{\Lambda}_t$  as a column vector, and  $\otimes$  is the Kronecker product. In Eq. 5,  $\circledast$  denotes circular convolution in image domain. In Eq. 6,  $\bar{\kappa}_t \in \mathbb{C}^{N^2 \times 1}$  is the image-domain kernel that captures the effects of frequency removal with  $\bar{\Lambda}_t$ ,  $x_0 = \mathcal{F}^{-1}(X_0)$  is the image sample at 0.

Based on linear systems theory, we note that the convolution operation in Eq. 6 can also be implemented via a matrix multiplication in image domain [76]:

$$x_t = D_t x_0, \quad (7)$$

where  $D_t \in \mathbb{C}^{N^2 \times N^2}$  is a doubly block circulant matrix (i.e., a special case of Toeplitz matrices). To construct  $D_t$  using the matrix form of  $\bar{\kappa}_t$ , i.e.,  $K \in \mathbb{C}^{N \times N}$ , a set of circulant matrices  $b_1, \dots, b_N \in \mathbb{C}^{N \times N}$  are created from individual rows of  $K$ :

$$b_i = \begin{bmatrix} K_{i,1} & K_{i,N} & \dots & K_{i,2} \\ K_{i,2} & K_{i,1} & \dots & K_{i,3} \\ \vdots & \vdots & \ddots & \vdots \\ K_{i,N} & K_{i,N-1} & \dots & K_{i,1} \end{bmatrix}, \quad (8)$$

where  $i \in [1, N]$  denotes row index, and  $K_{i,\cdot}$  is the  $i$ th row of  $K$ . The initial set of matrices are then organized in block

circulant form to obtain  $D_t$ :

$$D_t = \begin{bmatrix} b_1 & b_N & \dots & b_2 \\ b_2 & b_1 & \dots & b_3 \\ \vdots & \vdots & \ddots & \vdots \\ b_N & b_{N-1} & \dots & b_1 \end{bmatrix}. \quad (9)$$

**Recovery operator.** FDB's recovery operator  $G_\theta(x_t, t)$  produces an estimate of  $x_0$  given the current image sample  $x_t$  and step  $t$ . To derive the training objective for  $G_\theta(x_t, t)$ , we first observe that Eqs. 6 and 9 are equivalent. This reveals that the training objective of FDB admits a formal connection to the framework of generalized diffusion processes (GDP) [63], with the key distinction that the forward process of FDB given in Eq. 7 operates solely on stochastic degradations  $D_t$  driven by random k-space mask selection, whereas GDP combines linear degradations with Gaussian noise. For reference, the forward process for GDP is given by:

$$x_t = D_t x_0 + \sigma_t z, \quad (10)$$

where  $\sigma_t$  is a scalar that controls the noise variance. The score matching objective for the process in Eq. 10 is given as [77]:

$$L_{\text{GDP}} = \mathbb{E}_{t, x_t} [\|s_\theta(x_t, t) - \nabla_{x_t} \log(q(x_t))\|^2], \quad (11)$$

$$\equiv \mathbb{E}_{t, x_{(0,t)}} [\|s_\theta(x_t, t) - \nabla_{x_t} \log(q(x_t|x_0))\|^2], \quad (12)$$

where  $t \sim U(1, T_f)$ ,  $s_\theta(x_t, t)$  denotes the network parametrization of the score function, and the equivalence between Eqs. 11-12 up to a universal additive constant has been proven in the literature, including cases with  $\sigma_t = 0$  [63], [78]. Setting  $s_\theta(x_t, t) = (D_t G_\theta(x_t, t) - x_t)/\sigma_t^2$  and noting that  $q(x_t|x_0) = \mathcal{N}(D_t x_0, \sigma_t^2 I)$  as in [63], the objective in Eq. 12 becomes:

$$L_{\text{GDP}} = \mathbb{E}_{t, x_{(0,t)}} [\|D_t(G_\theta(x_t, t) - x_0)\|^2/\sigma_t^4]. \quad (13)$$

Note that for a noise-free degradation setting ( $\sigma_t = \epsilon$ , with a constant  $\epsilon \rightarrow 0$ ) and with  $D_t$  derived based on Eq. 9, Eq. 13 can be adapted to formulate FDB's training objective, while FDB's forward process remains as  $x_t = D_t x_0$  in Eq. 7:

$$L_{\text{FDB}} = \mathbb{E}_{t, x_{(0,t)}} [\|D_t(G_\theta(x_t, t) - x_0)\|^2/\epsilon^4]. \quad (14)$$

$$\equiv \mathbb{E}_{t, x_{(0,t)}} [\|D_t(G_\theta(x_t, t) - x_0)\|^2]. \quad (15)$$

Further noting that  $\bar{\Lambda}_t$  is diagonal with maximum eigenvalue 1, and  $\mathcal{F}^{-1}, \mathcal{F}$  are orthonormal transformations, we can make the following observation for the degradation operator:

$$\|D_t\| = \|\mathcal{F}^{-1} \bar{\Lambda}_t \mathcal{F}\| \leq \|\mathcal{F}^{-1}\| \|\bar{\Lambda}_t\| \|\mathcal{F}\| = 1. \quad (16)$$

Given that  $\|D_t(G_\theta(\cdot) - x_0)\| \leq \|D_t\| \|(G_\theta(\cdot) - x_0)\|$ , FDB can therefore be trained via a simplified upper-bound loss:

$$L_{\text{FDB}} = \mathbb{E}_{t, x_{(0,t)}} [\|(G_\theta(x_t, t) - x_0)\|^2]. \quad (17)$$

**A.2 Novel sampling algorithm:** Based on the theory of diffusion models, the discrete process in Eq. 10 can be expressed as the solution of the following stochastic differential equation (SDE) [77]:

$$dx_t = (\dot{D}_t \alpha) dt + \sqrt{\frac{d\sigma_t^2}{dt}} dw, \quad (18)$$

where  $\alpha$  denotes a single image sampled from  $q(x_0)$ , and  $w$  is a standard Wiener process. While this generalized form includes a noise-driven term,  $\sigma_t = 0$  in the case of FDB

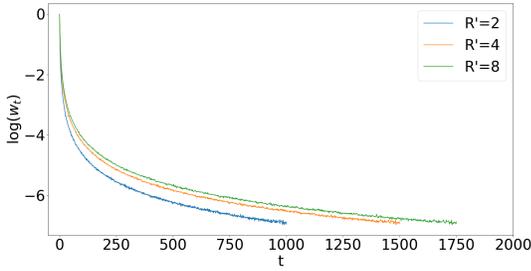


Fig. 4: The weighting parameter  $w_t$  for the correction term in Eq. 21 was learned as described in Eq. 32. Results shown for  $R'=2$  ( $T_f=1000$ ),  $R'=4$  ( $T_f=1500$ ), and  $R'=8$  ( $T_f=1750$ ) on the IXI dataset.

reduces the reverse SDE to:

$$dx_t = (\dot{D}_t \alpha) dt, \quad (19)$$

which can be solved numerically via the Euler-Maruyama method by approximating continuous derivatives as finite differences [44]:

$$x_{t-1} - x_t = (D_{t-1} - D_t)\alpha, \quad (20)$$

where  $D_0 = I$ . In general,  $\alpha$  will be unknown during reverse sampling, albeit an estimate of  $x_0$  can be obtained via the recovery operator as a substitute:  $\alpha = \{\tilde{x}_0(t) := G_\theta(x_t, t)\}$  [63]. Note that adopting the standard sampling procedure in Eq. 20 for FDB yields incremental imputation of missing frequency components between consecutive time steps. Here, we observe that this imputation performs hard dealiasing by only recovering  $n$  frequencies selected in the set  $S_t$ , without altering the value of previously recovered frequencies in the aggregated set  $S_{pre} = \bigcup_{\tau=T_f}^{t+1} S_\tau$ . Assuming that a frequency component  $k$  is imputed at step  $t_k$ , the recovered value is based on an imperfect estimate  $\tilde{x}_0(t_k) := G_\theta(x_{t_k}, t_k)$ , so it carries residual errors. Yet, based on Eq. 20, the value of  $k$  will not get updated at any other time step during  $t < t_k$  or  $t > t_k$ , so these residual errors cannot be corrected. In turn, the updates computed based on Eq. 20 cannot leverage the benefits of iterated Langevin sampling since they deviate from solving the true SDE that would be obtained when  $\alpha = x_0$ .

**Continually-corrected sampling.** To address this fundamental limitation, here we propose a novel continually-corrected sampling algorithm for FDB by incorporating a correction term that enables progressive dealiasing during reverse steps. In particular, we continue to check for residual errors on the previously imputed set of frequencies at each time step, and update their values to lower errors:

$$x_{t-1} - x_t = (D_{t-1} - D_t)\alpha + w_t D_t (\alpha - x_t), \quad (21)$$

where  $w_t \in [0, 1]$  is a weighting parameter, and the correction term effectively examines the difference  $(\tilde{x}_0(t) - x_t)$  to revise the values of frequency components imputed in previous steps  $\{T_f, \dots, t\}$ . In theory, assuming that  $\tilde{x}_0(t)$  is a good approximation to  $x_0$ , the correction term does not change the continuous-time SDE that is solved during reverse diffusion. This can be observed by setting  $\alpha = x_0$  for the correction term in Eq. 21 and noting that  $x_t = D_t x_0$ :

$$w_t D_t (\alpha - x_t) \approx w_t D_t (x_0 - x_t), \quad (22)$$

$$= w_t (D_t x_0 - D_t x_t), \quad (23)$$

$$= w_t (x_t - x_t) = 0. \quad (24)$$

In practice, the correction term can alter the original SDE to a degree depending on both the accuracy of the approximation  $\tilde{x}_0(t) \approx x_0$ , and the value of  $w_t$ . To understand the nature of alterations, we can reorganize the terms in Eq. 21, again using  $x_t = D_t x_0$  for simplification:

$$x_{t-1} - x_t = (D_{t-1} - D_t)\alpha + w_t D_t \alpha - w_t D_t x_t \quad (25)$$

$$x_{t-1} - (1 - w_t)x_t = (D_{t-1} - D_t(1 - w_t))\alpha. \quad (26)$$

When  $w_t \rightarrow 0$ , Eq. 25 converges to iterated Langevin sampling based on the original SDE formulation; and when  $w_t \rightarrow 1$ , it starts ignoring the image sample at step  $t$  to give instantaneous updates based on  $\alpha = \tilde{x}_0(t)$ . As we show in the next subsection, the optimal  $w_t$  schedule for FDB is attained by prescribing low  $w_t$  near step  $t = T_f$  (i.e., start-point), and high  $w_t$  near step  $t = 0$  (i.e., end-point). With this scheduling, the proposed algorithm performs iterated Langevin sampling near the start-point where  $\tilde{x}_0(t)$  is a poor estimate of  $x_0$ , deemphasizing the correction term. Contrarily, the algorithm transitions into performing instantaneous estimates solely based on  $\tilde{x}_0(t)$  near the end-point where it starts closely approximating  $x_0$ , hence emphasizing the correction term.

**Analytical derivation of the optimal  $w_t$ .** Compared to earlier iterations (i.e., near  $T_f$ ) where the set of recovered frequencies is compact, the proposed correction becomes more critical in later iterations as the set grows in size and the accuracy of  $\tilde{x}_0(t)$  increases. Thus, we adopted a learned schedule for  $w_t$  to enhance recovery. Observing that  $x_t = D_t x_0$ , Eq. 21 can be rearranged as:

$$x_{t-1} = (D_{t-1} - D_t)\alpha + D_t [w_t \alpha + (1 - w_t)x_t]. \quad (27)$$

The second term in Eq. 27 sets the values of previously recovered frequencies in  $x_{t-1}$  to a weighted average of  $\alpha = \tilde{x}_0(t)$  and  $x_t$ . Assuming  $\tilde{x}_0(t)$  reasonably approximates  $x_0$ ,  $w_t$  can be learned via a regression problem in Fourier domain:

$$\min_{w_t} \mathbb{E}_{t, X(0, t-1, t)} \underbrace{\|X_{t-1} - w_t X_0 - (1 - w_t)X_t\|_\Gamma^2}. \quad (28)$$

Since  $X_t^T X_0 = (\bar{\Lambda}_t X_0)^T (\bar{\Lambda}_t X_0) = \|X_t\|^2$ ,  $\Gamma$  is given as:

$$\Gamma = \|X_{t-1}\|^2 + w_t^2 \|X_0\|^2 + (1 - w_t)^2 \|X_t\|^2 - 2w_t \|X_{t-1}\|^2 - 2(1 - w_t) \|X_t\|^2 + 2w_t(1 - w_t) \|X_t\|^2, \quad (29)$$

$$= w_t^2 (\|X_0\|^2 - \|X_t\|^2) - 2w_t (\|X_{t-1}\|^2 - \|X_t\|^2) + (\|X_{t-1}\| - \|X_t\|)^2. \quad (30)$$

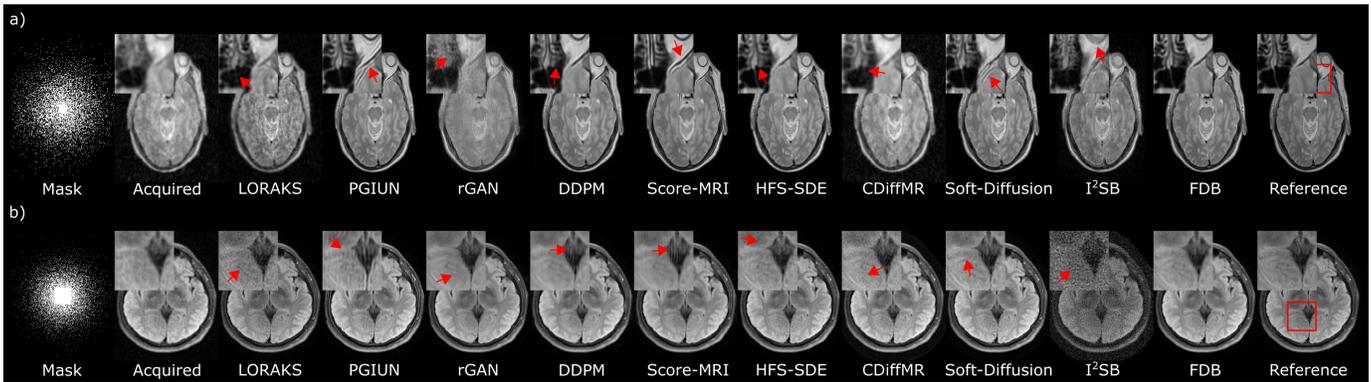
Note that the partial derivative of  $\Gamma$  with respect to  $w_t$  is:

$$\partial\Gamma/\partial w_t = 2w_t (\|X_0\|^2 - \|X_t\|^2) - 2 (\|X_{t-1}\|^2 - \|X_t\|^2). \quad (31)$$

Thus, the analytical solution to Eq. 28 can be expressed as:

$$w_t = \frac{\mathbb{E}_{t, X(t-1, t)} (\|X_{t-1}\|^2 - \|X_t\|^2)}{\mathbb{E}_{t, X(0, t)} (\|X_0\|^2 - \|X_t\|^2)}, \quad (32)$$

where the numerator and denominator reflect the expected energy differences between  $X_{t-1}$ ,  $X_t$  and  $X_0$ ,  $X_t$ , respectively, which can be estimated via a simple Monte-Carlo simulation over the training set. Note that  $w_t$  takes a maximum value of 1 attained at  $t = 1$ . Assuming that k-space energy distribution of MR images abide by an approximate inverse-power law (i.e.,  $\propto 1/r^\nu$ ),  $w_t$  should show exponential behavior [79]. Indeed, we find that  $w_t$  follows an exponential trend rising



**Fig. 5:** Within-domain reconstructions of representative cross-sections from (a) a PD acquisition in IXI at  $R = 10$  and (b) a FLAIR acquisition in fastMRI at  $R = 8$ . Images from competing methods are displayed along with the k-space undersampling mask, zero-filled Fourier reconstruction of undersampled data, and the reference image computed via Fourier reconstruction of fully-sampled data. Zoom-in windows are included to highlight differences among methods, and the red box on the reference image marks the zoomed region.

**TABLE I:** Within-domain performance on IXI at  $R = 4 - 10$  with 2D variable-density k-space sampling patterns. Results averaged across  $T_1$ ,  $T_2$ , PD contrasts. PSNR (dB) and SSIM (%) listed as mean $\pm$ std across the test set. Boldface marks the top performing method.

	$R = 4$		$R = 6$		$R = 8$		$R = 10$	
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
LORAKS	31.3 $\pm$ 1.7	63.2 $\pm$ 4.4	29.2 $\pm$ 1.6	57.4 $\pm$ 4.4	27.9 $\pm$ 1.3	55.7 $\pm$ 4.1	26.3 $\pm$ 1.3	46.7 $\pm$ 4.4
PGIUN	38.5 $\pm$ 2.4	96.9 $\pm$ 0.6	36.4 $\pm$ 2.4	97.0 $\pm$ 0.8	34.4 $\pm$ 2.4	96.7 $\pm$ 0.9	33.4 $\pm$ 2.5	96.0 $\pm$ 1.1
rGAN	36.6 $\pm$ 3.0	85.1 $\pm$ 7.1	33.9 $\pm$ 2.9	81.4 $\pm$ 7.2	32.6 $\pm$ 2.9	79.5 $\pm$ 7.3	30.7 $\pm$ 1.8	78.6 $\pm$ 8.6
DDPM	38.0 $\pm$ 2.5	97.6 $\pm$ 0.7	35.7 $\pm$ 2.5	96.7 $\pm$ 1.0	32.9 $\pm$ 2.6	95.4 $\pm$ 1.4	31.9 $\pm$ 2.6	94.5 $\pm$ 1.7
ScoreMRI	37.8 $\pm$ 2.5	97.6 $\pm$ 0.8	35.4 $\pm$ 2.5	96.7 $\pm$ 1.1	32.6 $\pm$ 2.5	95.6 $\pm$ 1.5	31.4 $\pm$ 2.6	94.8 $\pm$ 1.7
HFS-SDE	43.2 $\pm$ 2.8	<b>99.3<math>\pm</math>0.8</b>	40.1 $\pm$ 2.8	98.7 $\pm$ 0.9	36.3 $\pm$ 2.8	97.7 $\pm$ 1.1	35.1 $\pm$ 2.9	96.9 $\pm$ 1.2
CDiffMR	32.4 $\pm$ 2.4	86.0 $\pm$ 2.2	30.3 $\pm$ 2.3	82.4 $\pm$ 2.6	29.3 $\pm$ 2.4	80.7 $\pm$ 2.6	27.2 $\pm$ 2.3	75.4 $\pm$ 3.1
Soft-Diff.	38.5 $\pm$ 2.3	98.1 $\pm$ 0.6	36.0 $\pm$ 2.4	97.1 $\pm$ 0.9	33.2 $\pm$ 2.5	95.8 $\pm$ 1.3	32.0 $\pm$ 2.6	94.9 $\pm$ 1.7
I <sup>2</sup> SB	30.3 $\pm$ 2.6	83.0 $\pm$ 2.8	28.8 $\pm$ 2.6	80.6 $\pm$ 3.0	28.2 $\pm$ 2.6	79.9 $\pm$ 3.1	26.9 $\pm$ 2.6	76.3 $\pm$ 3.5
FDB	<b>43.7<math>\pm</math>2.9</b>	<b>99.3<math>\pm</math>0.9</b>	<b>40.8<math>\pm</math>2.8</b>	<b>98.8<math>\pm</math>0.9</b>	<b>36.9<math>\pm</math>2.6</b>	<b>97.8<math>\pm</math>1.0</b>	<b>36.1<math>\pm</math>2.7</b>	<b>97.3<math>\pm</math>1.3</b>

**TABLE II:** Within-domain performance on fastMRI at  $R = 4 - 10$  with 2D variable-density patterns. Results averaged across  $T_1$ ,  $T_2$ , FLAIR contrasts.

	$R = 4$		$R = 6$		$R = 8$		$R = 10$	
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
LORAKS	35.2 $\pm$ 2.4	92.0 $\pm$ 5.7	33.5 $\pm$ 2.3	90.5 $\pm$ 6.1	32.3 $\pm$ 2.3	89.5 $\pm$ 6.3	31.5 $\pm$ 2.3	88.6 $\pm$ 6.4
PGIUN	35.7 $\pm$ 2.3	93.7 $\pm$ 4.5	34.3 $\pm$ 2.2	91.8 $\pm$ 5.1	33.7 $\pm$ 2.2	90.9 $\pm$ 5.5	33.1 $\pm$ 2.3	90.0 $\pm$ 5.8
rGAN	35.9 $\pm$ 2.8	93.1 $\pm$ 4.7	33.9 $\pm$ 2.6	90.7 $\pm$ 5.5	32.9 $\pm$ 2.8	90.1 $\pm$ 5.9	31.1 $\pm$ 3.1	88.1 $\pm$ 6.4
DDPM	36.1 $\pm$ 2.3	93.7 $\pm$ 4.6	34.8 $\pm$ 2.3	92.2 $\pm$ 5.2	34.0 $\pm$ 2.3	91.1 $\pm$ 5.6	33.4 $\pm$ 2.2	90.1 $\pm$ 5.8
ScoreMRI	35.6 $\pm$ 2.2	93.4 $\pm$ 4.0	34.3 $\pm$ 2.2	91.8 $\pm$ 4.7	33.5 $\pm$ 2.2	90.5 $\pm$ 5.0	33.1 $\pm$ 2.2	89.9 $\pm$ 5.9
HFS-SDE	36.3 $\pm$ 2.1	92.9 $\pm$ 4.1	33.9 $\pm$ 1.8	89.0 $\pm$ 4.7	32.2 $\pm$ 1.7	85.1 $\pm$ 5.2	30.6 $\pm$ 1.8	81.1 $\pm$ 6.2
CDiffMR	31.4 $\pm$ 2.7	80.5 $\pm$ 7.4	31.2 $\pm$ 2.4	79.1 $\pm$ 7.2	31.2 $\pm$ 2.3	78.5 $\pm$ 8.1	31.1 $\pm$ 2.2	81.3 $\pm$ 7.1
Soft-Diff.	36.3 $\pm$ 2.6	91.1 $\pm$ 6.3	34.9 $\pm$ 2.3	89.9 $\pm$ 6.4	34.1 $\pm$ 2.1	89.2 $\pm$ 6.4	33.6 $\pm$ 2.0	88.7 $\pm$ 6.2
I <sup>2</sup> SB	28.4 $\pm$ 2.1	75.7 $\pm$ 4.5	26.9 $\pm$ 2.1	72.6 $\pm$ 4.8	26.2 $\pm$ 2.1	70.9 $\pm$ 5.0	25.7 $\pm$ 2.1	69.9 $\pm$ 5.0
FDB	<b>37.4<math>\pm</math>2.5</b>	<b>94.3<math>\pm</math>4.5</b>	<b>35.9<math>\pm</math>2.4</b>	<b>92.7<math>\pm</math>5.2</b>	<b>35.0<math>\pm</math>2.4</b>	<b>91.5<math>\pm</math>5.6</b>	<b>34.3<math>\pm</math>2.3</b>	<b>90.7<math>\pm</math>5.8</b>

from  $t = T_f$  to  $t = 1$ .

**A.3 MRI reconstruction with FDB:** Reconstruction is achieved by interleaving reverse diffusion steps for FDB with data-consistency projections. The trained prior maps an image corresponding to  $R'$ -fold undersampled data onto a high-quality MR image corresponding to fully-sampled data in  $T_f$  steps, by recovering  $n \approx N^2(R' - 1)/(R'T_f)$  frequency components in each step. Thus, to reconstruct an  $R$ -fold undersampled acquisition with  $N^2(R - 1)/R$  missing components, diffusion sampling must last for  $T_r = \lfloor T_f \frac{(R-1)R'}{(R'-1)R} \rfloor$  steps. As sinusoidal time encoding is used in FDB, sampling can be extrapolated to a broader time range (i.e.,  $T_r > T_f$ ) to account  $R > R'$  [80]. Yet,  $w_t$  is resampled across time to  $\tilde{w}_t$  to cover the range  $[w_1 w_{T_f}]$  in  $T_r$  steps. Sampling is initiated at  $T_r$  with the least-squares reconstruction of undersampled data,  $\hat{x}_{T_r} = A^H y$  [58]. To maintain stochasticity, here we

employ independent random instances of  $D_t$  during testing versus training. The mappings at step  $t$  are expressed as:

$$\hat{x}_{t-1} = \hat{x}_t + (D_{t-1} - D_t)\tilde{x}_0(t) + \tilde{w}_t D_t(\tilde{x}_0(t) - x_t), \quad (33)$$

$$\hat{x}_{t-1} = \hat{x}_{t-1} + A^H(y - A\hat{x}_{t-1}). \quad (34)$$

Finally,  $\hat{x}_0$  is taken as the reconstructed image.

## IV. METHODS

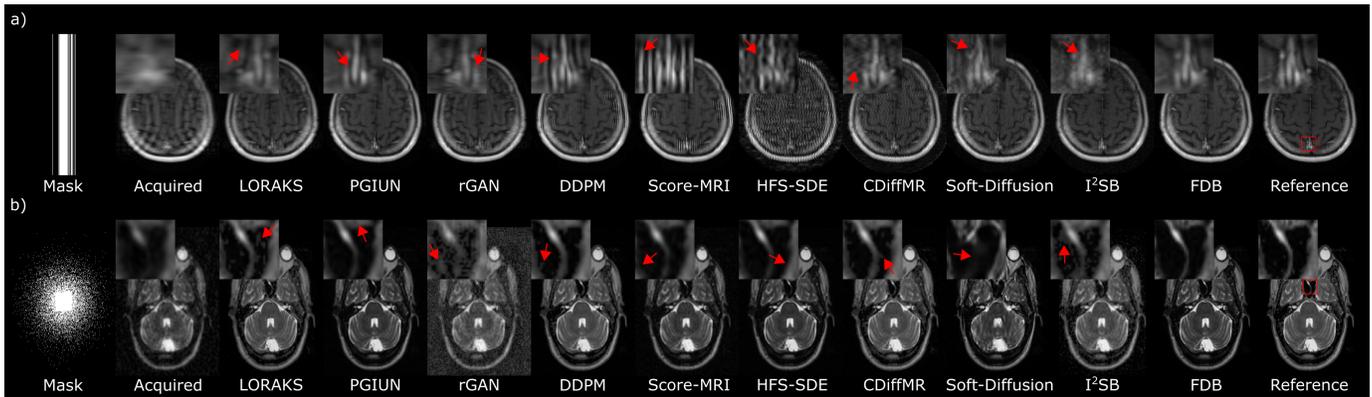
### A. Datasets

Experiments were performed on brain MRI data from IXI<sup>1</sup> and fastMRI [81]. Subjects were split into non-overlapping training, validation, test sets. In IXI, a subject-level split of (21,15,30) was used, and subjects were selected sequentially from the dataset. Coil-combined magnitude images for  $T_1$ ,  $T_2$  and PD contrasts were analyzed as single-coil acquisitions. For each contrast in a given subject, 100 axial cross-sections that contained brain tissue were selected, hence the training set comprised a total of 6300 cross-sections. In fastMRI, a subject-level split of (240,60,120) was used. Subjects were selected sequentially from the dataset, albeit only subjects whose volumes contained 10 cross-sections and more than 5 coil elements, and did not suffer from poor quality due to excessive noise were selected. Multi-coil complex k-space data for  $T_1$ ,  $T_2$  and FLAIR contrasts were analyzed. The training set comprised 7200 cross-sections. Geometric compression was used to derive 5 virtual coils that preserved over 90% of the data variance [82]. Retrospective undersampling at rates  $R = 4 - 10$  was used with normal density across the two-dimensional transverse plane [1]. To construct the imaging operators, coil sensitivities were estimated using ESPiRiT on a central calibration region [83]. Volumetric k-space data were inverse Fourier transformed and split across the readout dimension, and each cross-section was reconstructed individually.

### B. Competing Methods

FDB was demonstrated against nine state-of-the-art methods from the literature. Network models used two separate input and output channels to represent real and imaginary components. For each method, cross-validation procedures were used to select the values of key hyperparameters that maximized

<sup>1</sup><https://brain-development.org/ixi-dataset/>



**Fig. 6:** Cross-domain reconstructions of representative cross-sections. **(a)** A  $T_1$  acquisition in fastMRI at  $R = 8$  is shown. Models trained on 2D k-space sampling patterns were tested on 1D patterns under matched  $R$ . **(b)** A  $T_2$  acquisition in IXI at  $R = 8$  is shown. Models trained on multi-coil fastMRI data were tested on single-coil IXI data. Images from competing methods are displayed along with the k-space undersampling mask, zero-filled Fourier reconstruction of undersampled data, and the reference image computed via Fourier reconstruction of fully-sampled data. Zoom-in windows are included to highlight differences among methods, and the red box on the reference image marks the zoomed region.

**TABLE III:** Performance for models trained on 2D variable-density patterns, tested on 1D patterns for IXI at  $R = 4 - 8$ .

	$R = 4$		$R = 6$		$R = 8$	
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
LORAKS	24.3±0.9	42.0±2.6	22.5±0.7	37.0±2.5	21.3±0.7	34.2±2.7
PGIUN	32.0±2.6	95.9±1.3	29.6±2.5	91.6±2.6	27.0±2.4	89.1±2.9
rGAN	29.3±2.4	49.1±6.1	27.2±2.4	42.7±6.0	25.8±2.3	39.2±5.7
DDPM	30.1±2.8	94.3±1.9	26.7±2.7	89.6±3.0	25.1±2.5	85.5±3.6
ScoreMRI	29.8±2.7	94.8±1.8	26.5±2.6	90.8±2.9	25.0±2.5	87.0±3.5
HFS-SDE	33.1±3.0	96.4±1.4	28.6±3.0	91.8±2.6	25.5±2.9	85.7±4.0
CDiffMR	28.6±2.6	83.7±2.5	26.2±1.7	78.3±2.5	24.8±2.5	74.2±3.6
Soft-Diff.	30.5±2.6	94.5±2.0	27.8±2.6	89.1±2.7	25.2±2.6	85.6±4.0
I <sup>2</sup> SB	27.8±2.8	85.0±3.2	26.1±2.7	81.8±3.9	25.5±2.7	81.0±4.2
FDB	<b>34.0±2.9</b>	<b>96.7±1.2</b>	<b>30.4±2.8</b>	<b>93.5±2.0</b>	<b>28.1±2.7</b>	<b>89.8±2.7</b>

**TABLE IV:** Performance for models trained on 2D variable-density patterns, tested on 1D patterns for fastMRI at  $R = 4 - 8$ .

	$R = 4$		$R = 6$		$R = 8$	
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
LORAKS	28.1±2.3	85.3±5.6	25.9±2.3	80.5±6.2	24.5±2.2	76.9±6.5
PGIUN	30.1±2.0	89.4±5.5	27.9±1.7	82.3±6.0	26.4±1.7	80.6±6.2
rGAN	29.1±2.9	86.1±5.3	27.0±2.7	83.5±6.3	25.4±2.6	77.3±6.8
DDPM	29.4±2.0	88.8±5.5	27.4±1.5	84.1±5.6	25.7±1.5	80.6±6.5
ScoreMRI	29.0±1.8	88.1±5.0	26.7±1.6	83.5±5.7	25.2±1.5	79.7±6.2
HFS-SDE	24.3±1.2	61.1±4.9	22.6±1.1	52.6±4.9	21.6±1.1	47.9±4.9
CDiffMR	27.1±2.2	68.7±9.2	25.7±2.0	63.0±8.5	25.1±2.1	63.7±9.7
Soft-Diff.	29.9±1.7	84.8±5.3	28.2±1.5	79.8±5.4	26.9±1.5	75.9±5.6
I <sup>2</sup> SB	24.3±4.5	68.6±8.8	22.4±4.9	65.1±9.3	21.3±5.2	63.3±9.5
FDB	<b>30.6±2.2</b>	<b>89.5±6.0</b>	<b>29.0±2.0</b>	<b>85.9±6.9</b>	<b>27.8±2.0</b>	<b>82.7±7.3</b>

performance on the validation set. For FDB, these hyperparameters included the degradation levels at the start-point  $R'$ , and the number of epochs  $E$ . Modeling was performed via PyTorch on an Nvidia RTX 3090. Models were trained using the Adam optimizer with learning rate  $\eta = 10^{-4}$  and  $(\beta_1, \beta_2) = (0.5, 0.9)$ .

**FDB:** FDB used the architecture in [73].  $T_f=1000$ ,  $E=40$  epochs for IXI,  $E=10$  for fastMRI, and  $R'=2$  were used.

**LORAKS:** A low-rank reconstruction was considered [84]. The k-space neighborhood radius and the rank of the system matrix were set as (2,6) for IXI, and (2,30) for fastMRI.

**PGIUN:** A task-specific unrolled prior was considered [85].  $E=100$  was cross-validated.

**rGAN:** A task-specific GAN prior was considered [29]. Cross-validated hyperparameters were  $E=100$ , adversarial and pixel-wise loss weights of (1,100).

**DDPM:** A task-agnostic diffusion prior was implemented [47], [73]. Cross-validated hyperparameters were  $T=1000$ ,  $E=40$ .

**Score-MRI:** A score-based diffusion prior was implemented [50]. Cross-validated hyperparameters were  $T=1000$ ,  $E=40$ .

**HFS-SDE:** A score-based high frequency k-space diffusion prior was implemented [46]. Cross-validated hyperparameters were  $T=1000$ ,  $E=40$  for IXI,  $E=10$  for fastMRI.

**CDiffMR:** A cold diffusion prior was considered based on an asymptotic start-point derived via k-space undersampling [68]. Cross-validated hyperparameters were  $T=1000$ ,  $E=40$ .

**Soft-Diffusion:** A soft diffusion prior was considered based on an asymptotic start-point derived via k-space undersampling and additive Gaussian noise [63]. Cross-validated hyperparameters were  $T=1000$ ,  $E=40$ .

**I<sup>2</sup>SB:** A conventional diffusion bridge with start- and end-point identical to FDB was considered [69]. Cross-validated hyperparameters were  $T=200$ ,  $E=40$ .

### C. Analyses

For each dataset, model training was performed on data aggregated across multiple contrasts ( $T_1$ ,  $T_2$ , PD in IXI and  $T_1$ ,  $T_2$ , FLAIR in fastMRI). No explicit contrast information was provided during training, and data samples were randomly drawn for the aggregated training set. Task-specific models were trained for each  $R$  separately to maintain high performance. Meanwhile, task-agnostic models were trained without knowledge of  $R$ . Models were trained and tested on two-dimensional cross sections. Single-coil reconstructions were conducted on IXI, and multi-coil reconstructions were conducted on fastMRI. Reconstruction performance was assessed by quantifying PSNR and SSIM between recovered and reference images. Reference images were derived via Fourier reconstruction of fully-sampled acquisitions. Significance of performance differences among competing methods were assessed via non-parametric Wilcoxon signed-rank tests.

## V. RESULTS

### A. Comparison Studies

To improve its task relevance for MRI reconstruction, FDB builds a diffusion bridge that maps between realistically undersampled and fully-sampled data via stochastic degradation/recovery operators. We reasoned that the benefits of increased task relevance should be prominent in within-domain settings, where models are trained and tested on the

TABLE V:

Performance for models trained on 2D variable-density patterns, tested on 2D Poisson-disc patterns for IXI at $R = 4 - 8$ .	$R = 4$		$R = 6$		$R = 8$	
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
LORAKS	23.4±1.4	43.7±4.5	23.2±1.4	43.5±4.2	22.1±1.1	48.4±4.8
PGIUN	32.4±2.3	93.2±1.4	30.1±2.6	89.7±3.0	27.7±2.4	86.9±2.7
rGAN	29.3±2.2	47.0±5.9	26.9±2.1	42.8±5.1	25.4±2.2	38.3±5.2
DDPM	30.1±2.6	89.3±2.7	27.7±2.8	83.9±3.5	25.1±2.6	79.9±4.2
ScoreMRI	32.8±2.7	93.2±2.1	30.2±2.7	90.4±3.9	26.8±2.7	85.1±4.3
HFS-SDE	28.7±2.9	86.1±2.8	27.0±2.8	80.5±3.8	24.2±2.8	76.9±4.2
CDiffMR	23.7±2.4	67.0±3.9	23.5±2.6	65.7±3.9	22.9±2.5	64.1±4.2
Soft-Diff.	30.5±2.7	90.3±2.4	27.1±2.8	84.6±3.7	25.4±2.6	80.8±4.0
l <sup>2</sup> SB	25.4±2.7	72.7±4.2	24.9±2.7	72.4±4.3	24.6±2.7	72.2±4.7
FDB	<b>37.7±2.7</b>	<b>96.7±1.1</b>	<b>33.9±2.7</b>	<b>93.8±2.5</b>	<b>31.6±2.6</b>	<b>92.0±2.1</b>

TABLE VI:

Performance for models trained and tested on 1D patterns, for IXI at $R = 4 - 8$ .	$R = 4$		$R = 6$		$R = 8$	
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
LORAKS	24.3±0.9	42.0±2.6	22.5±0.7	37.0±2.5	21.3±0.7	34.2±2.7
PGIUN	32.2±2.5	96.0±1.4	29.7±2.8	92.0±2.1	27.2±2.0	89.1±2.5
rGAN	29.7±2.2	65.2±4.3	26.9±2.4	57.3±4.1	26.0±2.1	53.6±3.8
DDPM	30.1±2.8	94.3±1.9	27.4±2.6	88.7±3.0	25.1±2.5	85.5±3.6
ScoreMRI	29.8±2.7	94.8±1.8	26.5±2.6	90.2±2.9	25.0±2.5	87.0±3.5
HFS-SDE	33.8±3.2	<b>96.7±1.4</b>	30.6±2.8	92.1±2.8	26.8±2.9	87.0±3.7
CDiffMR	27.5±2.7	81.5±2.5	24.9±2.9	76.2±3.0	23.2±2.7	73.1±3.4
Soft-Diff.	30.2±2.4	93.8±1.9	26.5±2.5	87.9±3.8	24.8±2.8	84.9±3.7
l <sup>2</sup> SB	29.4±2.5	88.5±2.7	28.6±2.7	87.1±3.5	27.9±2.6	85.8±3.4
FDB	<b>34.0±2.9</b>	<b>96.7±1.2</b>	<b>31.2±2.7</b>	<b>92.4±2.4</b>	<b>28.1±2.7</b>	<b>89.8±2.7</b>

same dataset and sampling density. To assess this prediction, we compared FDB against a traditional method (LORAKS), task-specific priors (PGIUN, rGAN), task-agnostic diffusion priors (DDPM, Score-MRI, HFS-SDE), cold/soft diffusion priors (CDiffMR, Soft-Diffusion), and a conventional diffusion bridge (l<sup>2</sup>SB). Note that baseline models originally proposed with 1D k-space sampling (i.e., CDiffMR, HFS-SDE) were adapted for training with 2D sampling patterns to ensure consistency and systematic comparison across methods. Performance metrics in within-domain settings are listed in Table I for IXI, and in Table II for fastMRI. Overall, FDB is the top performer in all tasks and datasets ( $p < 0.05$ ), except for HFS-SDE that yields comparable SSIM at  $R=4$  on IXI similar to the within-domain settings. On average across tasks, FDB offers (PSNR, SSIM) improvements of (6.6dB, 22.4%) over the traditional method, (3.3dB, 5.3%) over task-specific priors, (2.6dB, 1.9%) over task-agnostic diffusion priors, (4.8dB, 8.5%) over cold/soft diffusion priors, (9.8dB, 19.2%) over the diffusion bridge, and (1.6dB, 2.7%) over the second-best method. All methods yield higher performance towards lower  $R$  values in both datasets. While the overall benefits that FDB provides over baselines follow a moderate inclination towards low  $R$  in IXI, the benefits are generally balanced across  $R$  in fastMRI, implying that spatial encoding from coil arrays can help maintain more consistent performance across  $R$ . Representative reconstructions are depicted in Fig. 5. LORAKS, l<sup>2</sup>SB and to a lesser degree CDiffMR, HFS-SDE show noise amplification; DDPM, Score-MRI, PGIUN show a degree of blur and ringing artifacts; rGAN, Soft-Diffusion can show visible noise and structural artifacts. In contrast, FDB yields images with more accurate depiction of tissue structure along with lower artifacts and noise. These results suggest that the increased task-relevance of the diffusion prior captured by FDB elicits superior performance in MRI reconstruction.

TABLE VII:

Cross-domain performance for models trained on multi-coil fastMRI data, and tested on single-coil IXI data at $R = 4 - 8$ .	$R = 4$		$R = 6$		$R = 8$	
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
LORAKS	31.3±1.7	63.2±4.4	29.2±1.6	57.4±4.4	27.9±1.3	55.7±4.1
PGIUN	34.2±2.7	92.9±1.6	32.8±2.7	92.2±1.6	31.9±2.8	92.1±1.8
rGAN	19.9±3.5	25.1±7.7	19.4±3.5	24.2±6.9	19.2±3.1	23.5±6.7
DDPM	34.8±2.7	93.6±1.4	33.3±2.6	92.4±1.6	32.2±2.6	92.0±1.8
ScoreMRI	31.8±2.7	86.0±2.6	30.2±2.7	83.4±2.9	29.4±2.8	82.3±3.1
HFS-SDE	34.8±2.6	93.6±1.4	33.3±2.6	92.4±1.6	32.1±2.5	92.0±1.7
CDiffMR	30.6±2.7	82.7±3.0	28.9±2.6	79.5±3.3	28.2±2.6	78.2±3.3
Soft-Diff.	23.7±2.2	60.5±5.4	19.7±2.6	53.3±5.4	16.7±2.7	45.3±5.5
l <sup>2</sup> SB	31.2±2.5	84.6±2.5	29.7±2.5	82.5±2.7	28.9±2.6	81.2±2.8
FDB	<b>36.0±2.7</b>	<b>95.4±1.0</b>	<b>34.1±2.6</b>	<b>93.8±1.2</b>	<b>32.7±2.7</b>	<b>93.0±1.3</b>

## B. Generalization

To ensure that the increased task relevance in FDB observed in within-domain settings does not come at the cost of generality, we further evaluated performance under alternative sampling strategies differing from the proposed design. Under matching undersampling rate  $R$ , we examined cases with (i) shifts from 2D variable-density to 1D sampling between training and test sets (Tables III, IV), (ii) shifts from 2D variable-density to 2D Poisson-disc sampling (Table V), and (iii) training and testing directly under 1D sampling (Table VI). FDB achieves the highest performance in all examined tasks ( $p < 0.05$ ), except for (iii) where HFS-SDE yields similar SSIM at  $R=4$ . On average, FDB offers (PSNR, SSIM) improvements of (7.8dB, 41.2%) over the traditional method, (3.0dB, 18.4%) over task-specific priors, (3.9dB, 6.5%) over task-agnostic diffusion priors, (4.9dB, 12.8%) over cold/soft diffusion priors, (5.7dB, 14.7%) over the diffusion bridge, and (2.0dB, 2.0%) over the second-best method. We also examined performance under domain shifts in dataset between training and test sets. Performance metrics under shifts from fastMRI to IXI are listed in Table VII. FDB achieves the highest performance in all examined tasks ( $p < 0.05$ ). On average, FDB offers (PSNR, SSIM) improvements of (4.8dB, 35.3%) over the traditional method, (8.0dB, 35.7%) over task-specific priors, (1.8dB, 4.3%) over task-agnostic diffusion priors, (9.6dB, 27.5%) over cold/soft diffusion priors, (4.3dB, 11.3%) over the diffusion bridge, and (0.8dB, 1.4%) over the second-best method.

All methods yield higher performance towards lower  $R$  values. Yet, in cases exploring alternative sampling strategies, FDB generally offers more prominent benefits over baselines towards higher  $R$  where the influence of the image prior used for reconstruction is expected to grow stronger. In case of shifts in dataset, the substantial improvements that FDB provides over task-specific and cold/soft diffusion priors, particularly in SSIM, are primarily driven by rGAN and Soft-Diffusion, respectively. We observed via inspection of ground-truth images in Fourier domain that fastMRI typically has more broad spread k-space spectra than IXI, hence images from the two datasets contain discrepant levels of high-frequency information. Poor performance of rGAN might then be attributed to the susceptibility of adversarial learning to domain shifts in the distribution of high-frequency information. Note that soft-diffusion priors are devised to perform joint denoising-dealiasing, attuning the relative emphasis on each transformation during reverse diffusion based on the energy levels of noise versus acquired frequency components in

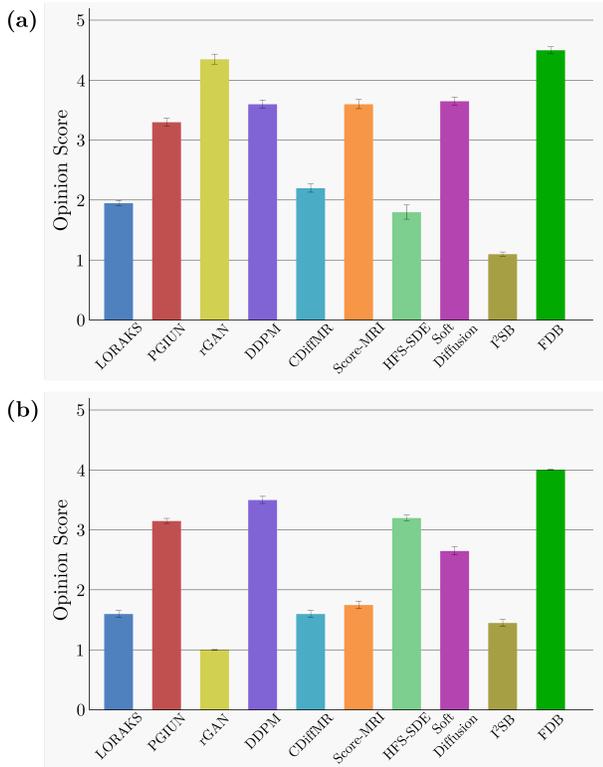


Fig. 7: Radiological opinion scores for competing methods in two representative tasks: (a) Within-domain reconstruction in fastMRI at  $R = 8$ , (b) Cross-domain reconstruction for fastMRI  $\rightarrow$  IXI at  $R = 8$ . Bar plots show mean  $\pm$  se of opinion scores across the test set.

TABLE VIII: Memory load (GB) and inference times (seconds) of competing methods for a cross-section in IXI at  $R = 4 - 10$ . Across  $R$ , memory load remains the same, whereas inference time can change.

	Mem.	Inf. Time			
		$R=4$	$R=6$	$R=8$	$R=10$
LORAKS	0.1	3.8	3.8	3.8	3.8
PGIUN	4.8	2.4	2.4	2.4	2.4
rGAN	2.6	0.1	0.1	0.1	0.1
DDPM	3.4	43.7	43.7	43.7	43.7
ScoreMRI	3.4	82.7	82.7	82.7	82.7
HFS-SDE	3.4	42.6	42.6	42.6	42.6
CDiffMR	3.7	36.2	39.9	42.1	44.0
Soft-Diff.	3.4	69.7	69.7	69.7	69.7
F <sup>2</sup> SB	10.4	154.0	154.0	154.0	154.0
FDB	4.4	47.1	51.8	54.2	55.9

undersampled data. Thus, poor performance of Soft-Diffusion might be attributed to the discrepant k-space spectra of the two datasets, causing ineffective noise suppression when the fastMRI-trained prior is tested on IXI, for which noise energy would be relatively more predominant towards the start-point. Collectively, these results suggest that FDB captures a more reliable prior across variations in sampling strategies and dataset distributions than competing methods, including diffusion priors.

Representative reconstructions are depicted in Fig. 6. Under shifts in sampling density, LORAKS, rGAN show residual aliasing artifacts; PGIUN shows structural inaccuracies; DDPM, CDiffMR, Score-MRI, HFS-SDE and to a lesser degree Soft-Diffusion show high-frequency ringing artifacts; F<sup>2</sup>SB shows high noise amplification. While competing methods are able to maintain relatively more reliable performance under shifts in dataset, LORAKS, rGAN, CDiffMR, F<sup>2</sup>SB suffer from noise amplification as particularly visible in background regions near tissue signals; DDPM, CDiffMR, HFS-

TABLE IX: Performance of reconstruction models when evaluated with various samplers. Columns show different model types, rows show different sampler types, and \* marks the native sampler associated with each model. Experiments were conducted on IXI at  $R = 8$ . For each reconstruction model (i.e., in each column), boldface marks the top performing sampler.

Sampler	DDPM		HFS-SDE		FDB	
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
EM (*DDPM)	<b>35.7<math>\pm</math>2.5</b>	96.7 $\pm$ 1.0	24.9 $\pm$ 1.7	78.2 $\pm$ 3.5	21.9 $\pm$ 1.7	74.8 $\pm$ 4.1
DPS	33.1 $\pm$ 2.0	<b>96.8<math>\pm</math>1.2</b>	26.3 $\pm$ 1.8	83.8 $\pm$ 3.3	22.7 $\pm$ 1.5	80.1 $\pm$ 4.2
PLMS	32.9 $\pm$ 1.9	95.9 $\pm$ 1.2	24.2 $\pm$ 1.7	75.1 $\pm$ 3.5	21.7 $\pm$ 1.7	71.4 $\pm$ 4.0
PC (*HFS-SDE)	33.0 $\pm$ 2.3	<b>96.8<math>\pm</math>1.0</b>	<b>36.3<math>\pm</math>2.8</b>	<b>97.7<math>\pm</math>1.1</b>	34.6 $\pm$ 2.1	96.7 $\pm$ 0.8
CCS (*FDB)	32.9 $\pm$ 2.3	<b>96.8<math>\pm</math>1.0</b>	26.9 $\pm$ 2.1	83.0 $\pm$ 4.0	<b>36.9<math>\pm</math>2.6</b>	<b>97.8<math>\pm</math>1.0</b>

SDE show a degree of spatial blur; PGIUN shows loss of detailed features; Soft-Diffusion can suffer from inaccurate depiction of anatomy. In comparison, FDB produces more accurate depiction of tissue structure, and relatively lower artifacts and noise than baselines. Note that task-specific priors that experience a training-test mismatch in sampling density or data distribution yield notable residual artifacts and noise amplification. While task-agnostic diffusion priors are relatively more resilient against shifts in dataset, they can suffer from suboptimal performance under challenging reconstruction tasks due to 1D sampling density. Meanwhile, FDB maintains a favorable trade-off between task-relevance and generalizability. These results suggest that the increased task-relevance of FDB does not compromise its superiority against baselines in cross-domain settings.

### C. Radiological Evaluation

While quantitative metrics such as PSNR and SSIM provide valuable objective benchmarks, they do not always align perfectly with visual image quality as perceived by radiologists. Subtle structural features may be visually more discernible despite only modest metric gains, whereas in other cases, substantial metric improvements may not translate into radiologically relevant differences. To capture this dimension, we conducted an observer study in which an expert neuro-radiologist assessed the similarity of images reconstructed from undersampled acquisitions to reference images derived from fully sampled acquisitions. Assessments were conducted using a 5-point Likert scale (1: lowest score, 5: highest score), while the observer was blinded to method names. Evaluations were performed for representative within-domain (fastMRI at  $R = 8$ ) and cross-domain (fastMRI  $\rightarrow$  IXI at  $R = 8$ ) reconstructions, on a set of 20 randomly selected test subjects per task. Radiological opinion scores of competing methods are displayed in Fig. 7. We find that FDB outperforms all competing methods across tasks ( $p < 0.05$ ). On average, FDB offers a score of 4.3, against scores of 1.8 for the traditional method, 3.0 for task-specific priors, 2.9 for task-agnostic diffusion priors, 2.5 for cold/soft diffusion priors, 1.3 for the conventional diffusion bridge. The relatively lower scores of competing methods were primarily driven by the presence of residual artifacts including spatial blurring and noise amplification in reconstructions. Taken together, these findings imply that improvements in quantitative performance metrics enabled by FDB are accompanied by radiologically-relevant enhancements in image quality.

**TABLE X:** Performance of FDB variants trained with  $R' = 2-16$  (degradation level during training), and tested for  $R = 4-8$  (undersampling rate) on IXI.

(a) FDB	$R = 4$		$R = 6$		$R = 8$	
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
$R' = 2$	<b>43.7±2.9</b>	<b>99.3±0.9</b>	40.8±2.8	<b>98.8±0.9</b>	<b>36.9±2.6</b>	<b>97.8±1.0</b>
$R' = 4$	42.5±2.9	99.1±0.7	38.6±2.3	98.0±0.5	35.0±2.7	96.7±1.0
$R' = 8$	42.6±2.8	99.2±0.5	38.3±2.2	98.0±0.5	34.6±2.5	96.5±1.0
$R' = 16$	42.4±3.1	99.1±0.6	38.2±1.9	98.0±0.5	34.6±2.6	96.6±1.0

(b) w/o adj.	$R = 4$		$R = 6$		$R = 8$	
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
$R' = 2$	42.3±2.9	99.2±0.8	40.7±3.0	98.7±0.6	35.7±2.8	97.4±1.0
$R' = 4$	42.5±2.9	99.1±0.7	<b>41.1±3.1</b>	98.6±0.6	36.5±2.8	97.6±0.9
$R' = 8$	39.6±2.7	98.5±0.7	38.2±2.4	97.7±0.9	34.6±2.5	96.5±1.0
$R' = 16$	37.6±2.5	97.8±0.7	35.3±1.9	96.3±0.9	32.5±2.4	95.0±1.1

### D. Computational Efficiency

Table VIII lists the memory load and inference time for competing methods. Among deep-learning models, the relatively compact rGAN has the lowest memory load, I<sup>2</sup>SB equipped with a relatively heavy backbone has the highest memory load, and remaining diffusion models including FDB have moderate demands comparable to each other. In terms of inference time, models that reconstruct images in a single forward pass require relatively short processing times. Meanwhile, diffusion models that employ iterated sampling algorithms naturally rely on longer computations. Among diffusion models, CDiffMR and FDB that alter their number of reverse diffusion steps based on the undersampling rate in the test set show modest increases in inference time towards higher  $R$ . Still, FDB has a moderate inference time compared against more demanding methods such as Score-MRI, Soft-Diffusion and I<sup>2</sup>SB.

### E. Influence of the Sampling Algorithm

A key element of FDB is its tailored sampling algorithm, Continually Corrected Sampling (CCS), which is integral to the proper functioning of its diffusion process. In conventional diffusion priors (e.g., DDPM, HFS-SDE), the reverse process operates via iterative image denoising: at each timestep, the network updates most or all of the image (or k-space) data, naturally refining previous estimates and self-correcting over time even with a standard sampling algorithm. In contrast, FDB's forward process removes non-overlapping subsets of k-space points at each step. Without CCS, previously recovered k-space points would remain fixed during the reverse process, depriving FDB of the iterative error-correction mechanism that is innate to conventional diffusion priors.

To investigate how the pairing between diffusion model and sampling algorithm influences performance, we conducted cross-pairing experiments by applying different sampling algorithms to trained DDPM, HFS-SDE, and FDB models. Specifically, we examined the vanilla Euler-Maruyama (EM) sampler [73] and CCS, along with sampling methods designed to improve diffusion priors through error correction, including Diffusion Posterior Sampling (DPS) [56], Pseudo Linear Multistep (PLMS) [86], and Predictor-Corrector (PC) [46]. Results at  $R = 8$  are presented in Table IX. In general, we find that each model typically achieves its best performance when paired with its native (i.e., originally proposed) sampling

**TABLE XI:** Performance of variants ablating method components on IXI for  $R = 4-8$ .

	$R = 4$		$R = 6$		$R = 8$	
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
FDB	<b>43.7±2.9</b>	<b>99.3±0.9</b>	<b>40.8±2.8</b>	98.8±0.9	<b>36.9±2.6</b>	97.8±1.0
w/o Fourier	17.0±2.8	50.9±5.2	16.7±1.2	50.7±3.1	17.1±2.6	51.4±4.9
w/o CCS	39.0±2.6	96.5±0.8	36.6±2.6	94.7±1.1	34.2±2.6	94.0±1.3
w/o dat.	30.5±1.6	93.3±1.7	28.9±1.5	91.0±2.3	28.0±1.5	89.6±2.6
w/o sto.	43.4±2.9	<b>99.3±0.8</b>	40.0±2.0	<b>99.0±0.3</b>	36.2±2.5	97.4±1.1
w/o rad.	43.0±2.7	<b>99.3±1.0</b>	40.0±2.0	<b>99.0±0.3</b>	36.8±2.5	<b>98.1±1.0</b>

**TABLE XII:** Performance of variants with different schedules for  $n$ ,  $w_t$  on IXI for  $R = 4-8$ .

		$R = 4$		$R = 6$		$R = 8$	
		PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
$n$	Linear	<b>43.7±2.9</b>	<b>99.3±0.9</b>	<b>40.8±2.8</b>	<b>98.8±0.9</b>	<b>36.9±2.6</b>	<b>97.8±1.0</b>
	Log	40.7±3.0	98.7±2.9	38.3±2.6	98.0±0.6	35.5±2.8	97.3±0.9
	Cosine	43.5±2.4	99.1±0.2	40.2±2.4	98.0±0.3	36.5±2.2	97.2±0.6
$w_t$	Learned	<b>43.7±2.9</b>	<b>99.3±0.9</b>	<b>40.8±2.8</b>	<b>98.8±0.9</b>	<b>36.9±2.6</b>	<b>97.8±1.0</b>
	Linear	24.4±2.4	85.1±3.1	22.6±1.2	82.6±2.6	20.8±2.4	80.2±2.7
	Exponential	25.4±1.1	88.5±2.0	24.5±1.1	87.7±2.0	24.1±1.2	88.3±1.5

algorithm. For instance, HFS-SDE performs optimally with the PC sampler, and FDB performs best with the CCS sampler. Furthermore, pairing models with non-native samplers typically causes performance degradations. For instance, HFS-SDE yields 36.3dB PSNR with the PC sampler, albeit its performance lowers to 26.9dB with the vanilla EM sampler. While very small SSIM improvements are occasionally seen for DDPM when using non-native samplers, they do not approach the performance achieved by FDB with CCS. These findings corroborate that CCS is an integral component of FDB that enables it to fully-realize the task-specific design of its diffusion process through continual error correction.

### F. Ablation Studies

A group of ablation studies were conducted to assess the influence of major method components and parameters in FDB. FDB is trained for a start-point of  $R'$ -fold undersampled data (corresponding to  $T_f$  timesteps). To improve flexibility and performance during inference, it adaptively adjusts the number of timesteps  $T_r$  while reconstructing  $R$ -fold undersampled test data. Table Xa lists reconstruction performance for  $R'=2-16$ . FDB models with lower  $R'$  values yield higher performance consistently across  $R$  prescribed for test acquisitions. We also examined the influence of adaptive adjustment of  $T_r$  by comparing FDB against a 'w/o adjust.' variant that always prescribed  $T_r = T_f$ . Table Xb lists performance of the variant model at  $R'=2-16$ . FDB achieves generally higher performance than 'w/o adjust.' for a given  $R'-R$  pair, and it maintains more consistent performance across increasing  $R'$  values at a given  $R$ . These findings suggest that performance improvements are viable in FDB by attuning  $R'$ , and that adaptive adjustment of  $T_r$  introduces a degree of robustness against  $R'-R$  mismatches.

Next, we examined the importance of method components related to the degradation operator and sampling algorithm by forming several FDB variants. A 'w/o Fourier' variant removed the Fourier constraint to express intermediate image samples as a weighted average of fully-sampled and under-sampled data as in conventional diffusion bridges [71]. A 'w/o

CCS' variant ablated the corrected sampling algorithm to use iterated Langevin sampling as in Eq. 20. A 'w/o dat.' variant ablated data-consistency projections during inference. A 'w/o sto.' variant ablated stochasticity by using shared instances of frequency removal matrices across training samples. A 'w/o rad.' variant removed the radius threshold in the degradation operator to promote unordered selection of removed frequency components across k-space. FDB generally yields superior performance against variants as listed in Table XI, except for 'w/o rad.' and 'w/o sto.' that occasionally yield on par or slightly higher SSIM. These results indicate that key method components that distinguish FDB from conventional diffusion bridges have important contributions to MRI reconstruction performance.

We also examined the importance of key parameters related to the schedules for the number of frequency components removed at each timestep ( $n$ ), and the weight of the correction term in the CCS algorithm ( $w_t$ ). For  $n$ , the linear schedule in FDB was compared against logarithmic and cosine schedules across timesteps. For  $w_t$ , the learned schedule in FDB was compared against linear and exponential schedules. As listed in Table XII, FDB outperforms variants in all tasks, indicating that the proposed parameter schedules contribute significantly to reconstruction performance.

## VI. DISCUSSION

FDB generates images via a dealiasing transformation that imputes unacquired k-space points in MRI data, eliciting two key differences in the context of diffusion modeling. First, FDB does not employ Gaussian noise addition that has been postulated to enable the training set to more uniformly cover the domain of the data distribution via a density smoothing effect in task-agnostic diffusion priors [44], [87]. Yet, corroborating recent reports based on analytical and heuristic approaches that showcase the feasibility of score-function learning via non-noise perturbations [62], [69], [72], here we demonstrate that FDB can effectively learn the distribution of fully-sampled MRI data without relying on noise addition. Second, maintaining stochasticity during reverse diffusion sampling has been suggested to aid correction of sampling errors [50], [58]. In the absence of noise addition, FDB maintains a degree of stochasticity via using random instances of frequency removal matrices across the diffusion process for each MR image, as we find that using fixed removal matrices does not yield notable benefits in MRI reconstruction.

Paralleling its performance in within-domain settings, FDB remains the top-performer among competing methods even under reasonable domain shifts in k-space sampling strategies and dataset. While task-agnostic diffusion priors rely on a noise-driven diffusion process untied to specific sampling patterns, granting them flexibility across sampling strategies, FDB enhances task relevance by incorporating 2D peripheral-to-central k-space undersampling in its diffusion process. Importantly, our experiments show that FDB maintains strong performance when training and testing across different sampling strategies (2D variable-density, 2D Poisson-disc, 1D) and undersampling rates ( $R'$  versus  $R$ ), indicating that design

elements promoting task specificity do not unduly compromise generalization to unseen or significantly different patterns. That said, clinical confirmation of reliability against substantial domain shifts in sampling patterns and anatomical content remains an important direction for future work. Recent studies have reported that adaptation of image priors to individual subjects can boost generalization to atypical anatomy [38], [41]. During inference, adaptation requires optimization of network parameters that can be computationally burdening for diffusion priors with thousands of sampling steps, so accelerated sampling methods would be required to improve practicality [45]. The preprocessing steps in IXI may alter the characteristics of MRI data, potentially reducing the complexity of reconstruction tasks and yielding an over-optimistic view on acceleration rates achievable in clinical use. Note, however, that all of the main experiments in the current study were also conducted on the fastMRI dataset that contains unprocessed multi-coil complex data, and the results from the two datasets provide a convergent view corroborating the superiority of FDB over competing methods.

Several technical limitations can be addressed to help enhance the performance and utility of FDB. Diffusion models typically employ Langevin sampling across thousands of steps, resulting in characteristically slow inference. Unlike task-agnostic diffusion priors, FDB initiates sampling with the least-squares reconstruction of undersampled acquisitions, recently shown to improve efficiency [58]. Yet, FDB has longer inference times compared to adversarial priors with one-step image sampling. To improve efficiency, several acceleration strategies can be adopted including adversarial learning to enable reverse diffusion over large step sizes [45], interleaved sampling followed by a refinement procedure [52], and distillation of trained priors [88]. Future studies are warranted to assess the influence of acceleration on FDB's performance.

The observed improvements in reconstruction performance with the CCS algorithm suggest that estimation accuracy for  $\tilde{x}_0(t)$  is relatively poor near the start-point of the diffusion process, hampering estimation accuracy and computational efficiency. A further approach to improve sampling accuracy could be to adopt recursive estimation of  $\tilde{x}_0(t)$  at each time step [74], which could shorten the number of steps required to approach convergence. FDB's recovery operator predicts the clean image at 0 given the image sample at  $t$ , and then degrades it back to obtain the sample at  $t - 1$ . To mitigate potential deviations in acquired k-space components from their originally measured values, FDB hence uses data-consistency projections between consecutive reverse diffusion steps. Improved efficiency might be attained by adopting data-consistent sampling algorithms that avoid these projections altogether [70]. FDB constructs its diffusion process via undersampling degradations for MRI reconstruction, yet this task-specialization might hamper performance in non-reconstruction tasks (e.g., deblurring). To build a multi-task model while avoiding performance losses, training sets with mixed samples could be used derived from the different degradation operators associated with each task. A task identifier could also be employed to allow task-specific modulation of feature maps across the recovery network [89].

Another area for improvement concerns learning strategies in FDB. Here, key hyperparameters including the level of degradation at the start-point, and the number of training epochs were tuned via cross-validation. Since the tuned set of parameters is compact and modest performance changes were observed under moderate drifts from optimal values, we reason that practical implementations will not face major computational burden. When needed, hyperparameter tuning might be performed on smaller training sets to enhance practicality. Learning on training sets with mixed MRI contrasts might be improved by employing contrast-specific modulation of feature maps [89], by training contrast-specific priors [15], or by jointly reconstructing across multiple contrasts [29], [90]. In cases where acquisition of fully-sampled data to permit supervised learning is challenging [91], [92], reliance on fully-sampled data can be lowered via self-supervised, cycle consistent, or low-rank assisted learning [48], [87], [93], [94].

Other areas for improvement include the network architecture, degradation operator, and domain of operation for FDB. To boost contextual sensitivity without elevating computational burden, hybrid architectures comprising convolutional and transformer blocks might be utilized in FDB [30], [39], [95]. Here, performance reconstructions were obtained via a degradation operator following peripheral-to-central order for frequency removal in compliance with variable-density sampling patterns. If desired, the degradation operator can be adopted to implement Poisson or uniform patterns [83]. It remains future work to assess the ideal sampling patterns for FDB. Here, FDB was implemented as an image-domain diffusion model. While a k-space implementation could be expected to offer similar benefits [51], [53], it remains future work to assess the relative performances of variants in image-domain versus k-space. Dual-domain implementations might enable performance improvements at the expense of elevated model complexity [20].

## VII. CONCLUSION

In this study, we presented a novel diffusion bridge, FDB, for accelerated MRI reconstruction. During forward diffusion, FDB's stochastic degradation operator performs random spatial-frequency removal to map directly between fully-sampled and undersampled data. For image reconstruction, FDB performs reverse diffusion with a novel continually-corrected sampling algorithm to progressively dealias MR images. Demonstrations on brain MRI indicate that FDB outperforms state-of-the-art reconstruction methods based on non-diffusion and diffusion priors.

## REFERENCES

- [1] M. Lustig, D. Donoho, and J. M. Pauly, "Sparse mri: The application of compressed sensing for rapid mr imaging," *Magn. Reson. Med.*, 58, 6, 1182–1195, 2007.
- [2] X. Zhang *et al.*, "Image reconstruction with low-rankness and self-consistency of k-space data in parallel mri," *Med. Image Anal.*, 63, 101687, 2020.
- [3] D. Liang, J. Cheng, Z. Ke, and L. Ying, "Deep magnetic resonance image reconstruction: Inverse problems meet neural networks," *IEEE Signal Process. Mag.*, 37, 1, 141–151, 2020.

- [4] C. M. Sandino, J. Y. Cheng, F. Chen, M. Mardani, J. M. Pauly, and S. S. Vasanawala, "Compressed sensing: From research to clinical practice with deep neural networks: Shortening scan times for magnetic resonance imaging," *IEEE Sig. Process. Mag.*, 37, 1, 117–127, 2020.
- [5] G. Zeng *et al.*, "A review on deep learning mri reconstruction without fully sampled k-space," *BMC Med. Imaging*, 21, 1, 195, 2021.
- [6] K. Hammernik *et al.*, "Physics-driven deep learning for computational magnetic resonance imaging," *arXiv:2203.12215*, 2022.
- [7] Y. Chen *et al.*, "Ai-based reconstruction for fast mri—a systematic review and meta-analysis," *Proc. IEEE*, 110, 2, 224–245, 2022.
- [8] F. Lam, X. Peng, and Z.-P. Liang, "High-dimensional mr spatiospectral imaging by integrating physics-based modeling and data-driven machine learning: Current progress and future directions," *IEEE Signal Process. Mag.*, 40, 2, 101–115, 2023.
- [9] S. Wang *et al.*, "Knowledge-driven deep learning for fast MR imaging: Undersampled MR image reconstruction from supervised to unsupervised learning," *Magn. Reson. Med.*, 92, 2, 496–518, 2024.
- [10] —, "Accelerating magnetic resonance imaging via deep learning," *IEEE ISBI*, 2016, 514–517.
- [11] J. Schlemper, J. Caballero, J. V. Hajnal, A. Price, and D. Rueckert, "A Deep Cascade of Convolutional Neural Networks for MR Image Reconstruction," *IPMI*, 2017, 647–658.
- [12] K. Hammernik *et al.*, "Learning a variational network for reconstruction of accelerated MRI data," *Magn. Reson. Med.*, 79, 6, 3055–3071, 2017.
- [13] B. Zhu, J. Z. Liu, B. R. Rosen, and M. S. Rosen, "Image reconstruction by domain transform manifold learning," *Nature*, 555, 7697, 487–492, 2018.
- [14] H. K. Aggarwal, M. P. Mani, and M. Jacob, "MoDL: Model-Based deep learning architecture for inverse problems," *IEEE Trans. Med. Imaging*, 38, 2, 394–405, 2019.
- [15] S. U. H. Dar, M. Özbey, A. B. Çatli, and T. Çukur, "A transfer-learning approach for accelerated MRI using deep neural networks," *Magn. Reson. Med.*, 84, 2, 663–685, 2020.
- [16] K. Kwon, D. Kim, and H. Park, "A parallel MR imaging method using multilayer perceptron," *Med. Phys.*, 44, 12, 6209–6224, 2017.
- [17] C. M. Hyun, H. P. Kim, S. M. Lee, S. Lee, and J. K. Seo, "Deep learning for undersampled MRI reconstruction," *Phys. Med. Biol.*, 63, 13, 135007, 2018.
- [18] T. M. Quan, T. Nguyen-Duc, and W.-K. Jeong, "Compressed sensing MRI reconstruction with cyclic loss in generative adversarial networks," *IEEE Trans. Med. Imaging*, 37, 6, 1488–1497, 2018.
- [19] G. Yang *et al.*, "DAGAN: Deep de-aliasing generative adversarial networks for fast compressed sensing MRI reconstruction," *IEEE Trans. Med. Imaging*, 37, 6, 1310–1321, 2018.
- [20] T. Eo, Y. Jun, T. Kim, J. Jang, H.-J. Lee, and D. Hwang, "KIKI-net: cross-domain convolutional neural networks for reconstructing undersampled magnetic resonance images," *Magn. Reson. Med.*, 80, 5, 2188–2201, 2018.
- [21] J. Yoon *et al.*, "Quantitative susceptibility mapping using deep neural network: QSMnet," *NeuroImage*, 179, 199–206, 2018.
- [22] M. Mardani *et al.*, "Deep generative adversarial neural networks for compressive sensing MRI," *IEEE Trans. Med. Imaging*, 38, 1, 167–179, 2019.
- [23] C.-M. Feng, Y. Yan, H. Fu, L. Chen, and Y. Xu, "Task transformer network for joint mri reconstruction and super-resolution," *MICCAI*, 2021, 307–317.
- [24] Z. Ramzi, C. G R, J.-L. Starck, and P. Ciuciu, "NC-PDNet: A Density-Compensated Unrolled Network for 2D and 3D Non-Cartesian MRI Reconstruction," *IEEE Trans. Med. Imaging*, 41, 7, 1625–1638, 2022.
- [25] G. Wang, T. Luo, J.-F. Nielsen, D. C. Noll, and J. A. Fessler, "B-spline parameterized joint optimization of reconstruction and k-space trajectories (bjork) for accelerated 2d mri," *IEEE Trans. Med. Imaging*, 41, 9, 2318–2330, 2022.
- [26] J. Lyu *et al.*, "Region-focused multi-view transformer-based generative adversarial network for cardiac cine mri reconstruction," *Med. Image Anal.*, 85, 102760, 2023.
- [27] W. Gan *et al.*, "Self-supervised deep equilibrium models with theoretical guarantees and applications to mri reconstruction," *IEEE Trans. Comput. Imaging*, 9, 796–807, 2023.
- [28] Q. Liu, Q. Yang, H. Cheng, S. Wang, M. Zhang, and D. Liang, "Highly undersampled magnetic resonance imaging reconstruction using autoencoding priors," *Magn. Reson. Med.*, 83, 1, 322–336, 2020.
- [29] S. U. Dar, M. Yurt, M. Shahdloo, M. E. Ildiz, B. Tinaz, and T. Çukur, "Prior-guided image reconstruction for accelerated multi-contrast mri via generative adversarial networks," *IEEE J. Sel. Top. Signal Process.*, 14, 6, 1072–1087, 2020.

- [30] P. Guo, Y. Mei, J. Zhou, S. Jiang, and V. M. Patel, "Reconformer: Accelerated mri reconstruction using recurrent transformer," *arXiv:2201.09376*, 2022.
- [31] F. Liu, A. Samsonov, L. Chen, R. Kijowski, and L. Feng, "Santis: Sampling-augmented neural network with incoherent structure for MR image reconstruction," *Magn. Reson. Med.*, 82, 5, 1890–1904, 2019.
- [32] B. Yaman, S. A. H. Hosseini, and M. Akcakaya, "Zero-shot physics-guided deep learning for subject-specific mri reconstruction," *NeurIPS Workshop DLIP*, 2021.
- [33] Z. Wang *et al.*, "One-dimensional deep low-rank and sparse network for accelerated MRI," *IEEE Trans. Med. Imaging*, 42, 1, 79–90, 2023.
- [34] A. Pramanik, S. Bhave, S. Sajib, S. D. Sharma, and M. Jacob, "Adapting model-based deep learning to multiple acquisition conditions: Adamodl," *Magn. Reson. Med.*, 2023.
- [35] X. Zhao, T. Yang, B. Li, and X. Zhang, "Swingun: A dual-domain swin transformer-based generative adversarial network for mri reconstruction," *Comput. Biol. Med.*, 153, 106513, 2023.
- [36] K. C. Tezcan, C. F. Baumgartner, R. Luechinger, K. P. Pruessmann, and E. Konukoglu, "Mr image reconstruction using deep density priors," *IEEE Trans. Med. Imaging*, 38, 7, 1633–1642, 2018.
- [37] U. S. Kamilov, C. A. Bouman, G. T. Buzzard, and B. Wohlberg, "Plug-and-play methods for integrating physical and learned models in computational imaging: Theory, algorithms, and applications," *IEEE Signal Process. Mag.*, 40, 1, 85–97, 2023.
- [38] D. Narnhofer, K. Hammernik, F. Knoll, and T. Pock, "Inverse gans for accelerated mri reconstruction," *Wavelets and Sparsity XVIII*, 11138, SPIE, 2019, 381–392.
- [39] Y. Korkmaz, S. U. Dar, M. Yurt, M. Özbey, and T. Cukur, "Unsupervised mri reconstruction via zero-shot learned adversarial transformers," *IEEE Trans. Med. Imaging*, 41, 7, 1747–1763, 2022.
- [40] G. Luo, N. Zhao, W. Jiang, E. S. Hui, and P. Cao, "Mri reconstruction using deep bayesian estimation," *Magn. Reson. Med.*, 84, 4, 2246–2261, 2020.
- [41] M. Z. Darestani and R. Heckel, "Accelerated mri with un-trained neural networks," *IEEE Trans. Comput. Imaging*, 7, 724–733, 2021.
- [42] G. Elmas *et al.*, "Federated learning of generative image priors for mri reconstruction," *IEEE Trans. Med. Imaging*, 42, 7, 1996–2009, 2023.
- [43] K. P. Slavkova *et al.*, "An untrained deep learning method for reconstructing dynamic mr images from accelerated model-based data," *Magn. Reson. Med.*, 89, 4, 1617–1633, 2023.
- [44] Y. Song, L. Shen, L. Xing, and S. Ermon, "Solving inverse problems in medical imaging with score-based generative models," *ICLR*, 2022.
- [45] A. Güngör *et al.*, "Adaptive diffusion priors for accelerated MRI reconstruction," *Med. Image Anal.*, 88, 102872, 2023.
- [46] C. Cao *et al.*, "High-frequency space diffusion model for accelerated mri," *IEEE Trans. Med. Imaging*, 43, 5, 1853–1865, 2024.
- [47] G. Luo, M. Blumenthal, M. Heide, and M. Uecker, "Bayesian mri reconstruction with joint uncertainty estimation using diffusion models," *Magn. Reson. Med.*, 90, 1, 295–311, 2023.
- [48] H. Peng *et al.*, "One-shot generative prior in hankel-k-space for parallel imaging reconstruction," *IEEE Trans. Med. Imaging*, 1–1, 2023.
- [49] A. Jalal, M. Arvinte, G. Daras, E. Price, A. G. Dimakis, and J. Tamir, "Robust compressed sensing mri with deep generative priors," *NeurIPS*, 34, 2021, 14938–14954.
- [50] H. Chung and J. C. Ye, "Score-based diffusion models for accelerated mri," *Med. Image Anal.*, 80, 102479, 2022.
- [51] Y. Xie and Q. Li, "Measurement-conditioned denoising diffusion probabilistic model for under-sampled medical image reconstruction," *arXiv:2203.03623*, 2022.
- [52] C. Peng, P. Guo, S. K. Zhou, V. Patel, and R. Chellappa, "Towards performant and reliable undersampled mr reconstruction via diffusion model sampling," *MICCAI*, 2022, 623–633.
- [53] Z. Xiao *et al.*, "Diffusion model based on generalized map for accelerated mri," *NMR Biomed.*, 37, 12, e5232, 2024.
- [54] B. Levac, S. Kumar, A. Jalal, and J. I. Tamir, "Accelerated motion correction with deep generative diffusion models," *Magn. Reson. Med.*, 92, 2, 853–868, 2024.
- [55] W. Bian, A. Jang, L. Zhang, X. Yang, Z. Stewart, and F. Liu, "Diffusion Modeling with Domain-conditioned Prior Guidance for Accelerated MRI and qMRI Reconstruction," *IEEE Trans. Med. Imaging*, 1–1, 2024.
- [56] H. Chung, J. Kim, M. McCann, M. Klasky, and J. Ye, "Diffusion posterior sampling for general noisy inverse problems," *ICLR*, 2023.
- [57] H. Chung, S. Lee, and J. Ye, "Decomposed diffusion sampler for accelerating largescale inverse problems," *ICLR*, 2024.
- [58] H. Chung, B. Sim, and J. C. Ye, "Come-closer-diffuse-faster: Accelerating conditional diffusion models for inverse problems through stochastic contraction," *CVPR*, 2022, 12413–12422.
- [59] Z.-X. Cui *et al.*, "Spirit-diffusion: Self-consistency driven diffusion model for accelerated mri," *IEEE Trans. Med. Imaging*, 2024.
- [60] B. Ozturkler, M. Mardani, A. Vahdat, J. Kautz, and J. M. Pauly, "Regularization by denoising diffusion process for MRI reconstruction," *NeurIPS Workshop DLIP*, 2023.
- [61] Z.-X. Cui *et al.*, "Physics-Informed DeepMRI: k-Space Interpolation Meets Heat Diffusion," *IEEE Trans. Med. Imaging*, 43, 10, 3503–3520, 2024.
- [62] A. Bansal *et al.*, "Cold diffusion: Inverting arbitrary image transforms without noise," *NeurIPS*, 2023, 41259–41282.
- [63] G. Daras, M. Delbracio, H. Talebi, A. G. Dimakis, and P. Milanfar, "Soft diffusion: Score matching for general corruptions," *Trans. Mach. Learn. Res.*, 2023.
- [64] S. Lee, H. Chung, J. Kim, and J. C. Ye, "Progressive deblurring of diffusion models for coarse-to-fine image synthesis," *arXiv:2207.11192*, 2022.
- [65] A. Radmanesh *et al.*, "Exploring the Acceleration Limits of Deep Learning Variational Network-based Two-dimensional Brain MRI," *Rad. Artif. Intell.*, 4, 6, e210313, 2022.
- [66] Y. Chen *et al.*, "AI-Based Reconstruction for Fast MRI—A Systematic Review and Meta-Analysis," *Proc. IEEE*, 110, 2, 224–245, 2022.
- [67] Y. Balaji *et al.*, "ediffi: Text-to-image diffusion models with an ensemble of expert denoisers," *arXiv:2211.01324*, 2022.
- [68] J. Huang, A. I. Aviles-Rivero, C.-B. Schönlieb, and G. Yang, "Cdiffr: Can we replace the gaussian noise with k-space undersampling for fast mri?" *MICCAI*, 2023, 3–12.
- [69] G.-H. Liu, A. Vahdat, D.-A. Huang, E. A. Theodorou, W. Nie, and A. Anandkumar, "I<sup>2</sup>SB: Image-to-Image Schrödinger Bridge," *ICML*, 2023, 22042–22062.
- [70] Z. Fabian, B. Tinaz, and M. Soltanolkotabi, "Diracdiffusion: Denoising and incremental reconstruction with assured data-consistency," *arXiv:2303.14353*, 2023.
- [71] M. Delbracio and P. Milanfar, "Inversion by direct iteration: An alternative to denoising diffusion for image restoration," *Trans. Mach. Learn. Res.*, 2023.
- [72] H. Chung, J. Kim, and J. C. Ye, "Direct diffusion bridge using data consistency for inverse problems," *arXiv:2305.19809*, 2023.
- [73] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," *NeurIPS*, 33, 6840–6851, 2020.
- [74] F. Arslan, B. Kabas, O. Dalmaz, M. Ozbey, and T. Çukur, "Self-consistent recursive diffusion bridge for medical image translation," *arXiv:2405.06789*, 2024.
- [75] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," *NeurIPS*, 33, 2020, 6840–6851.
- [76] A. Jain, *Fundamentals of Digital Image Processing*. Prentice Hall, 1989.
- [77] Y. Song, L. Shen, L. Xing, and S. Ermon, "Solving inverse problems in medical imaging with score-based generative models," *arXiv:2111.08005*, 2021.
- [78] P. Vincent, "A connection between score matching and denoising autoencoders," *Neural Comput.*, 23, 7, 1661–1674, 2011.
- [79] D. L. Ruderman and W. Bialek, "Statistics of natural images: Scaling in the woods," *Phys. Rev. Lett.*, 73, 814–817, 1994.
- [80] A. Q. Nichol and P. Dhariwal, "Improved denoising diffusion probabilistic models," *ICML*, 2021, 8162–8171.
- [81] F. Knoll *et al.*, "fastMRI: A publicly available raw k-space and DICOM dataset of knee images for accelerated MR image reconstruction using machine learning," *Rad. Artif. Intell.*, 2, 1, e190007, 2020.
- [82] T. Zhang, J. M. Pauly, S. S. Vasanawala, and M. Lustig, "Coil compression for accelerated imaging with cartesian sampling," *Magn. Reson. Med.*, 69, 2, 571–582, 2013.
- [83] M. Uecker *et al.*, "Esprit—an eigenvalue approach to autocalibrating parallel mri: where sense meets grappa," *Magn. Reson. Med.*, 71, 3, 990–1001, 2014.
- [84] J. P. Haldar and J. Zhuo, "P-loraks: low-rank modeling of local k-space neighborhoods with parallel imaging data," *Magn. Reson. Med.*, 75, 4, 1499–1514, 2016.
- [85] J. Jiang, Z. He, Y. Quan, J. Wu, and J. Zheng, "PGIUN: Physics-Guided Implicit Unrolling Network for Accelerated MRI," *IEEE Trans. Comput. Imaging*, 2024.
- [86] L. Liu, Y. Ren, Z. Lin, and Z. Zhao, "Pseudo numerical methods for diffusion models on manifolds," 2022.
- [87] A. Aali, M. Arvinte, S. Kumar, and J. I. Tamir, "Solving inverse problems with score-based generative priors learned from noisy data," *arXiv:2305.01166*, 2023.
- [88] H. A. Bedel and T. Çukur, "DreaMR: Diffusion-driven counterfactual explanation for functional MRI," *arXiv:2307.09547*, 2023.

- 
- [89] O. Dalmaz *et al.*, "One model to unite them all: Personalized federated learning of multi-contrast mri synthesis," *Med. Image Anal.*, 94, 103121, 2024.
  - [90] L. Xiang *et al.*, "Deep-learning-based multi-modal fusion for fast MR reconstruction," *IEEE Trans. Biomed. Eng.*, 66, 7, 2105–2114, 2019.
  - [91] B. Yaman, S. A. H. Hosseini, S. Moeller, J. Ellermann, K. Uğurbil, and M. Akçakaya, "Self-supervised learning of physics-guided reconstruction neural networks without fully sampled reference data," *Magn. Reson. Med.*, 84, 6, 3172–3191, 2020.
  - [92] J. Liu, Y. Sun, C. Eldeniz, W. Gan, H. An, and U. S. Kamilov, "Rare: Image reconstruction using deep priors learned without groundtruth," *IEEE J. Sel. Top. Signal Process.*, 14, 6, 1088–1099, 2020.
  - [93] Z.-X. Cui *et al.*, "Self-Score: Self-Supervised Learning on Score-Based Models for MRI Reconstruction," *arXiv:2209.00835*, 2022.
  - [94] T. Xiang, M. Yurt, A. B. Syed, K. Setsompop, and A. Chaudhari, "Ddm<sup>2</sup>: Self-supervised diffusion mri denoising with generative diffusion models," *arXiv preprint arXiv:2302.03018*, 2023.
  - [95] C.-M. Feng *et al.*, "Multimodal transformer for accelerated mr imaging," *IEEE Trans. Med. Imaging*, 42, 10, 2804–2816, 2023.